Title

Multimodal Sensor Fusion and Machine Learning for Layer-Wise Monitoring of Laser Powder Bed Fusion

ISBN

9798190926132

Author

Satterlee, Nicholas

Publication Date

2026-09-11

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO
SAN DIEGO STATE UNIVERSITY

Multimodal Sensor Fusion and Machine Learning for Layer-Wise Monitoring of Laser Powder Bed Fusion

A Dissertation submitted in partial satisfaction of the requirements
for the degree Doctor of Philosophy

in

Engineering Sciences (Mechanical and Aerospace Engineering)

by

Nicholas A. Satterlee

Committee in charge:

University of California San Diego
Professor Olivia Greaves, Co-Chair
Professor Boris Kramer
Professor Michael Todd

San Diego State University
Professor John Kang, Co-Chair
Professor Eugene Olevsky

2026

The Dissertation of Nicholas A. Satterlee is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

Co-Chair

Co-Chair

University of California San Diego
San Diego State University

2026

DEDICATION

In memory of my brother, Austen, and his lasting love of science.

&

To my beautiful wife, Brigid, for her love, positivity, and encouragement.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to acknowledge Professor John Kang for his support as the co-chair of my committee. His continuous guidance through some of the most challenging times of my life facilitated a large part of my success with this project.

Chapter 3 in part contains material as it appears in Parallel multi-layer sensor fusion for pipe leak detection using multi-sensors and machine learning. N. Satterlee, X. Zuo, C.-W. Lee, C.-W. Park, and J. S. Kang, Engineering Applications of Artificial Intelligence, vol. 153, p. 110923, 2025. The dissertation author was the primary investigator and author of this paper.

Chapter 3 in part contains material as it appears in Feature-informed machine learning for detecting material deformation and failure in aluminum pipes under bending load using acoustic emission sensors. X. Zuo, N. Satterlee, C.-W. Lee, I.-G. Choi, C.-W. Park, and J. S. Kang, *Materials & Design*, vol. 254, p. 114087, 2025. The dissertation author developed feature-informed network architecture and performed analysis on the results.

VITA

2012	Bachelor of Science, University of California, Berkeley
2012-2013	Physics Intern, Lawrence Livermore National Laboratory
2013-2014	Co-founder/CEO, Seed LLC
2014-2018	Nuclear Fuel Engineer, Southern Nuclear
2020-2021	Teaching Associate, San Diego State University
2021	Master of Science, San Diego State University
2026	Teaching Associate, San Diego State University
2026	Doctor of Philosophy in Engineering Sciences (Mechanical and Aerospace Engineering), University of California San Diego and San Diego State University

PUBLICATIONS

T. Grippi, E. Torresani, R. Jiang, A. Maximenko, D. Giuntini, M. v. Maris, J. Kang, **N. Satterlee**, R. Bordia, E. Olevsky, "Continuum Meso-Macro Multiscale Modeling of Anisotropy During Sintering of Parts Produced Via Additive Manufacturing", *npj Advanced Manufacturing* (submitted)

J. Boomgaarden, **N. Satterlee**, R. H. Schmitt, K. Geiger, J. S. Kang, "Smart Metrology Techniques for Additive Manufacturing," *International Journal of Advanced Manufacturing Technology*, 144, 261–289, 2026.

X. Zuo, **N. Satterlee**, R. Guwalani, C Park and J. S. Kang, "Comparative Analysis of Machine Learning Approaches for Leak Detection in Pipelines Using Hydrophone Acoustic Sensing", *Measurement*, vol. 261, p. 120005, 2026.

N. Satterlee, R. Jiang, E. Torresani, T. Grippi, A. Maximenko, M.P.F.H.L. van Maris, D. Giuntini, E. Olevsky, J. S. Kang, "Machine Learning-Driven Characterization of Anisotropic Microstructures in Powder-Based Additive Manufactured Metal Parts", *Progress in Additive Manufacturing*, 2026.

N. Satterlee, E. Torresani, R. Jiang, M. P. F. H. L. van Maris, D. Giuntini, E. Olevsky, J. S. Kang, "Machine Learning-Enabled Characterization of 3D Particle Contacts from 2D Microstructural Images in Powder-Based Metal Additive Manufacturing", *TMS 2026 155th Annual Meeting & Exhibition Supplemental Proceedings, Cham: Springer Nature Switzerland*, 2026, pp. 569–583.

N. Satterlee, X. Zuo, K. S. Moon, S. Q. Lee, M. Peterson, J. S. Kang*, "Sentence-Level Silent Speech Recognition Using a Wearable EMG/EEG Sensor System with AI-Driven Sensor Fusion and Language Model", *Sensors*, Vol. 25, Iss. 19, p. 6168, 2025.

X. Zuo, **N. Satterlee**, C. Lee, I. Choi, C. Park* and J. S. Kang*, "Feature-Informed Machine Learning for Detecting Material Deformation and Failure in Aluminum Pipes Under Bending Load Using Acoustic Emission Sensors", *Materials & Design*, Vol. 254, pp. 114087, 2025.

N. Satterlee, X. Zuo, C. Lee, C. Park and J. S. Kang, "Parallel Multi-Layer Sensor Fusion for Pipe Leak Detection Using Multi-Sensors and Machine Learning", *Engineering Applications of Artificial Intelligence*, Vol. 153, p. 110923, 2025.

J. S. Kang, K. S. Moon, S. Q. Lee, **N. Satterlee**, “A Wearable Silent Text Input System Using EMG and Piezoelectric Sensors”, *Sensors*, Vol. 25, Iss. 8, pp. 2642, 2025.

K. S. Moon, J. S. Kang, S. Q Lee, J. Thompson, **N. Satterlee**, “Wireless Mouth Motion Recognition System Based on EEG-EMG Sensors for Severe Speech Impairments”, *Sensors*, Vol. 24, Iss. 13, pp. 4125, 2024.

N. Satterlee, R. Jiang, E. Olevsky, E. Torresani, X. Zuo, and J. S. Kang, “Robust Image-Based Cross-Sectional Grain Boundary Detection and Characterization using Machine Learning”, *Journal of Intelligent Manufacturing*, Vol. 36, pp. 3067-3095, 2024.

N. Satterlee, E. Torresani, E. Olevsky, J. S. Kang, “Automatic Detection and Characterization of Porosities in Cross-Section Images of Metal Parts Produced by Binder Jetting using Machine Learning and Image Augmentation”, *Journal of Intelligent Manufacturing*, Vol. 35, pp. 1281-1303, 2024.

N. Satterlee, E. Torresani, E. Olevsky, J. S. Kang, “Comparison of Supervised Machine Learning Algorithms for Automatic Detection of Porosities in Powder-Based Additive Manufactured Metal Parts”, *The International Journal of Advanced Manufacturing Technology,* Vol. 120, pp. 6761–6776, 2022.

Y.P. Lin, R.M. Fawcett, S. S. Desilva, D.R. Lutz, M.O. Yilmaz, P. Davis, R.A. Rand, P.E. Cantonwine, R.B. Rebak, R. Dunavant, **N. Satterlee**, “Path Towards Industrialization of Enhanced Accident Tolerant Fuel”, *Top Fuel 2018*, Prague, 2018, A0141.

Descalle, M.-A. and Brown, D. and Beck, B. and Summers, C. and Vranas, P. and Barnowski, R. and **Satterlee, N**., “Verification and Validation of LLNL’s Evaluated Nuclear Data Library (ENDL2011)’, *2013 International Nuclear Data Conference for Science and Technology*, March 2013, pg. 55-56.

ABSTRACT OF THE DISSERTATION

Multimodal Sensor Fusion and Machine Learning for Layer-Wise Monitoring of Laser Powder Bed Fusion

by

Nicholas A. Satterlee

Doctor of Philosophy in Engineering Sciences (Mechanical and Aerospace Engineering)

University of California San Diego, 2026
San Diego State University, 2026

Professor Olivia Greaves, Co-Chair
Professor John Kang, Co-Chair

Laser powder bed fusion (LPBF) enables fabrication of complex metal components, but its coupled multiphysics drives multiple defect-formation mechanisms. Porosity is a critical defect because it forms through transient, spatially localized process events and can degrade mechanical reliability. Localized in-situ pore detection is therefore essential for resolving conditions associated with individual pore formation, advancing understanding of LPBF

process–structure relationships, and supporting more reliable process monitoring and qualification.

This dissertation develops a Multimodal Sensor Fusion Machine Learning (MSFML) framework for localized LPBF pore prediction by integrating heterogeneous sensor data with process- and physics-informed knowledge. The proposed Parallel Multi-Layer Multimodal Sensor Fusion (PMMSF) method progressively fuses one-dimensional sensor signals, image data, photometric-stereo products, process parameters, physics-derived features, and multi-layer context.

The experimental system integrated three thermal-emission channels, acoustic emission, laser-triggered acquisition, coaxial melt-pool imaging, and off-axis photometric stereo imaging during the fabrication of nominally dense coupons under varying laser power and scan speed conditions. A comprehensive data preparation workflow was developed to synchronize and register these heterogeneous measurements. Post-build X-ray computed tomography (XCT) was used to establish three-dimensional pore ground truth and link individual pores to their corresponding in-situ process histories.

Single-sensor models, multimodal sensor combinations, fusion strategies, process-informed weighting, and physics-informed learning were evaluated. The full multimodal framework with process-informed weighting and physics-informed features achieved the best performance, with 95% test accuracy, 97% AUC, and an F1 score of 87%. F1 improved by 10 percentage points over the best single-sensor model, 5 points over the best two-sensor model, and 3 points over the best three-sensor model. Process-informed weighting and physics-informed features each provided an additional 2-point F1 improvement.

In summary, this dissertation establishes a framework that advances localized in-situ pore prediction by integrating multimodal sensing with process- and physics-informed machine learning. The framework enables investigation of transient pore-formation signatures and LPBF process–structure relationships while providing a generalizable approach for sensor-based monitoring of other complex manufacturing processes.

Chapter 1 Introduction

1.1 Background and Motivation

Additive manufacturing (AM) offers greater design flexibility and reduced material waste compared to traditional manufacturing techniques. Among metal AM technologies, laser powder bed fusion (LPBF) is one of the most widely adopted because high-resolution metal components with complex geometries, short lead times, and a wide range of materials can be produced. In LPBF, a focused laser selectively melts regions of a powder bed according to a programmed scan path. After each layer is scanned, a new powder layer is spread, and the process is repeated until the part is completed. This layer-wise manufacturing method enables complex part geometries that are difficult or impossible to fabricate using conventional manufacturing processes (Frazier; King et al.; DebRoy et al.).

Despite these advantages, LPBF still faces major challenges in ensuring consistent part quality. The process involves rapid melting, solidification, vaporization, powder redistribution, melt-pool flow, and repeated thermal cycling. These coupled physical phenomena can produce pores, lack-of-fusion defects, surface defects, cracks, balling, dimensional errors, and microstructural variability. These defects are especially important for critical applications because they can reduce fatigue life, increase scatter in mechanical properties, and increase the cost and complexity of qualification. The transition of LPBF from prototyping to critical-part production is limited by confidence in part quality, and there remains a need to connect AM process conditions to structure and properties (King et al.; DebRoy et al.).

Porosity is the primary defect considered in this dissertation. Pores are important because they can act as stress concentrators and reduce mechanical reliability. They are also difficult to monitor because they are spatially localized and can form during short-lived process events. A pore may form within a small region of a single melt track, may be influenced by the previous-layer surface condition, and may not be visible at the top surface after the layer is completed. Post-process characterization methods such as X-ray computed tomography (XCT), cross-section imaging, Archimedes density, and mechanical testing can characterize the final part, but they do not provide real-time information during the build (King et al.; DebRoy et al.; Everton et al.; Grasso and Colosimo).

The work in this dissertation emphasizes porosity in nominally dense LPBF parts. Keyhole-related pore formation is an important motivation because high-energy-density processing can generate a deep vapor depression, or keyhole, in the melt pool. If the keyhole becomes unstable, it can fluctuate, collapse, and form bubbles that are trapped by the solidification front. Keyhole fluctuation, collapse, and moving keyhole-tip instability have been shown to be directly connected to pore formation through synchrotron X-ray imaging (Huang et al.; Zhao et al.). However, the source of an individual stochastic pore is not always known from XCT data alone. A pore may originate from keyhole instability, local powder-bed irregularity, entrapped gas, lack of fusion, part-edge effects, or interactions between multiple process variables. Therefore, it is not assumed in this dissertation that every detected pore is a keyhole pore. Instead, XCT-detected porosity in dense LPBF parts is studied while keyhole-mode physics is used as one important motivation for sensor selection and physics-informed feature design.

In-situ monitoring provides a path toward real-time pore detection, but each sensor measures only an indirect signature of the process. Pyrometers measure thermal emission from the melt-pool region. A coaxial camera captures melt-pool geometry and intensity. Acoustic emission (AE) measures high-frequency dynamic events associated with rapid material or melt-pool activity. Off-axis images and photometric-stereo heightmaps capture powder-bed and layer-surface features that may act as defect precursors. Process data provides laser position, laser trigger, layer number, and nominal process parameters. These sensors are complementary, but they are heterogeneous in sampling rate, dimensionality, noise, spatial resolution, and physical interpretation (Everton et al.; Grasso and Colosimo; Gorgannejad et al.).

In-situ monitoring data have been used for LPBF pore or defect prediction in prior studies. Thermographic in-situ monitoring has been combined with deep learning to predict porosity-related defects, and multimodal process monitoring has been combined with X-ray radiography for localized keyhole-pore prediction during LPBF (Oster et al.; Gorgannejad et al.). These studies demonstrate that in-situ sensor signals can contain defect-relevant information. However, localized pore prediction in a practical LPBF monitoring system remains challenging because in-situ signals must be synchronized with the laser path and registered to three-dimensional ground-truth pore locations.

The central challenge is that in-situ sensor signals must be connected to actual defect locations. XCT provides three-dimensional pore ground truth after the build, but the XCT volume is initially in its own coordinate system. To use XCT labels for machine learning, the XCT pore locations must be segmented, converted to physical units, registered to the machine coordinate system, and synchronized with the in-situ sensor data. Without this registration, a sensor window may be labeled incorrectly as pore or non-pore, resulting in label noise and unreliable model interpretation (Oster et al.; Gorgannejad et al.; Hocker et al.).

The overall goal of this dissertation is to develop a robust and accurate multimodal sensor-fusion framework for near-real-time LPBF pore monitoring. The framework integrates multi-sensor data acquisition, laser-trigger-based synchronization, machine-coordinate mapping, part-boundary segmentation, XCT pore segmentation, XCT-to-process registration, localized query-volume labeling, physics-informed feature extraction, process-parameter-informed fusion, and PMMSF.

1.2 Research Problem

The research problem addressed in this dissertation is the localized prediction of pore defects in LPBF using the proposed pore prediction framework.

This problem is difficult for several reasons. First, pores are sparse and may be stochastic. A printed part contains a very large number of possible spatial locations, but only a small fraction of those locations contains pores. Second, the defect source is not always known. Although keyhole-mode physics motivates part of this study, detected pores may also be influenced by powder-bed defects, local surface condition, entrapped gas, lack-of-fusion mechanisms, part-edge effects, or interactions between multiple process variables. Third, in-situ sensors do not directly measure pores. Instead, indirect process signatures such as thermal emission, melt-pool shape, acoustic transients, and layer-surface condition are measured. Fourth, the sensors are heterogeneous. The AE and pyrometer signals are one-dimensional time series, the coaxial camera produces high-speed images, the off-axis system produces layer-wise images and heightmaps, and the process data are tabular scan-path data. Fifth, the XCT labels are collected after printing and must be spatially registered to the process coordinate system before they can be used for supervised learning (Everton et al.; Grasso and Colosimo; Oster et al.; Gorgannejad et al.). Finally, localized pore prediction requires a fusion strategy that can combine heterogeneous sensor streams without collapsing useful modality-specific information. Conventional approaches often rely on manual feature extraction, classical classifiers, input concatenation, channel fusion, or final-layer head fusion (Gorgannejad et al.; Tang et al). These methods can be effective, but they may compress sensor-specific information too early, allow a dominant modality to suppress weaker modalities, or delay cross-modal interaction until after important intermediate features have already been lost.

Defect prediction has been demonstrated in prior studies using single sensors, large analysis windows, implanted voids, synchrotron imaging, or macroscopic labels. These studies are valuable, but the problem of localized pore prediction in a practical LPBF monitoring system is not fully solved. A deployable monitoring system requires sensors that can be integrated with the printer, a synchronization workflow that maps sensor signatures to laser position, a ground-truth method that localizes real pores in three dimensions, and a fusion model that can combine heterogeneous sensor information without allowing one dominant sensor to suppress weaker but potentially useful modalities (Oster et al.; Gorgannejad et al.; Satterlee et al.; Baltrusaitis et al.).

Accordingly, synchronization, XCT registration, and label generation are critically important preprocessing steps. The quality of the model depends directly on the quality of the labels. If the XCT pore locations are shifted relative to the process coordinates, then positive and negative labels are assigned incorrectly.

1.3 Research Hypothesis

The main hypothesis of this dissertation is that localized pore prediction in LPBF can be improved by fusing multi-layer multi-modal sensor data through a physics- and process-informed Parallel Multi-Layer Multimodal Sensor Fusion architecture.

This hypothesis is based on three assumptions. First, pore-relevant process signatures are distributed across multiple sensor modalities. Thermal signals, melt-pool geometry, acoustic activity, powder-bed condition, and process parameters each provide partial information about pore formation. Second, multi-layer context may contain useful information because local surface condition, remelting, and thermal history can influence pore formation in the current or subsequent layers. Third, physics- and process-informed features can guide the fusion model by providing information about operating regime and melt-pool stability.

1.4 Research Objectives

The overall goal of this dissertation is to develop a robust and accurate multimodal Sensor Fusion Machine Learning (SFML) framework for near-real-time LPBF pore monitoring. To achieve this goal, three objectives are identified.

The first objective is to develop a synchronized and spatially registered multimodal dataset for localized LPBF pore prediction. A multi-sensor monitoring system is integrated and operated on an LPBF machine using AE, three pyrometer channels, coaxial imaging, off-axis imaging, photometric-stereo heightmaps, and process data. These data are collected during builds with varied laser power and scan speed. Because the sensor streams are acquired at different sampling rates, coordinate systems, and temporal resolutions, a spatiotemporal registration and labeling workflow is developed to assign machine coordinates to the heterogeneous sensor data and align them with XCT-derived pore ground truth. Pores are segmented from XCT volumes, and their morphology and locations are registered to the machine-coordinate system. The registered pore geometry is then used to generate localized pore/no-pore labels by determining the overlap between each process-centered sensor query volume and the XCT pore geometry, establishing direct correspondence between in-situ sensor signatures and three-dimensional pore defects for machine learning.

The second objective is to develop and evaluate the PMMSF method for multimodal LPBF pore prediction. PMMSF is designed to combine heterogeneous sensor data while preserving modality-specific feature extraction and enabling progressive interaction between sensor features across multiple network layers. The framework consists of sensor-specific

Feature Extraction Networks, a Sensor Fusion Network, and a Network Head that progressively fuse multimodal features throughout the network. The method is evaluated through comparisons of individual sensors, reduced sensor combinations, the full sensor suite, input concatenation, head fusion, and PMMSF. These comparisons are used to determine which sensing modalities and fusion strategies provide the strongest pore-prediction performance and the best performance-to-complexity tradeoff.

The third objective is to develop and evaluate physics- and process-informed learning within PMMSF. Process parameters and physics-derived features are incorporated into the framework to determine whether domain knowledge provides additional predictive information beyond the multimodal sensor data alone. Process-informed inputs include laser power, scan speed, hatch spacing, spot size, layer height, volumetric energy density, and areal energy density. Physics-informed inputs include melt-pool-depth proxy features, melt-pool-depth derivatives, relative temperature proxies, melt-pool geometry features, and topographic features. Process-informed PMMSF and physics-informed PMMSF are compared with the baseline PMMSF configuration to quantify the contribution of each form of domain information.

1.5 Scope of the Dissertation

This dissertation focuses on localized porosity prediction in nominally dense 316L stainless steel parts produced by laser powder bed fusion (LPBF). The study is limited to a single material processed on a single LPBF machine.

Pore ground truth is obtained using XCT, and the capabilities of the XCT system define the scope of this work. The XCT data have a voxel size of approximately 7.6 µm; therefore, pores smaller than approximately 20 µm in diameter cannot be reliably identified and are excluded from this study.

The pore population is not assumed to arise from a single formation mechanism. Instead, pores may result from keyhole instability, powder-bed irregularities, gas entrapment, lack of fusion, or complex interactions among multiple process variables. Accordingly, the proposed framework is designed to predict pore occurrence without explicitly identifying the underlying formation mechanism.

This dissertation also investigates multimodal sensor fusion strategies. While head fusion provides a stronger baseline than direct input concatenation by extracting modality-specific features before fusion, it performs cross-modal interaction only at the final network stage after each modality has been compressed into a latent representation. Consequently, it may fail to capture informative intermediate interactions among thermal, acoustic, optical, and topographic features. To address this limitation, the proposed Parallel Multi-Layer Multimodal Sensor Fusion (PMMSF) framework performs progressive, layer-wise fusion, enabling earlier and richer cross-modal feature interactions.

Microstructure prediction, mechanical property prediction, and direct physics-based modeling of pore formation mechanisms, such as keyhole collapse dynamics, are beyond the scope of this dissertation and are considered future research directions.

1.6 Dissertation Organization

Chapter 2 reviews LPBF pore formation, in-situ monitoring, XCT ground truth, sensor synchronization, machine learning for AM defect prediction, multimodal fusion, and physics-informed learning. Chapter 3 presents the feature-engineering and sensor-fusion approach used to convert synchronized multimodal data and registered XCT pore morphology into localized pore/no-pore predictions. Chapter 4 describes the LPBF platform, multimodal acquisition system, sensor suite, process matrix, data-collection software, and raw data products. Chapter 5 describes data preparation, including process-coordinate mapping, part segmentation, AE and coaxial synchronization, off-axis image registration, XCT pore segmentation, XCT-to-process registration, and localized query-volume label generation. Chapter 6 reports model results, including dataset statistics, individual-sensor performance, sensor-combination studies, fusion-architecture comparisons, process-informed fusion, and physics-informed fusion. Chapter 7 discusses the implications, limitations, uncertainty sources, and deployment relevance of the results. Chapter 8 summarizes the dissertation conclusions and future research directions.

Chapter 2 Background

2.1 Overview

Laser powder bed fusion (LPBF) is a coupled thermal, fluid, optical, mechanical, and metallurgical process. During printing, powder is spread across the build plate and selectively melted by a focused laser as shown in Figure 2.1. The resulting melt pool undergoes rapid heating, melting, vaporization, fluid flow, solidification, and reheating as subsequent layers are deposited. Because the final part quality depends on local interactions between process parameters, powder condition, scan strategy, layer history, and melt-pool dynamics, defect formation cannot be described only by the nominal laser power or scan speed. Instead, defect formation must be interpreted as a local and transient process (King et al.; Debroy et al.; Khairallah et al.).

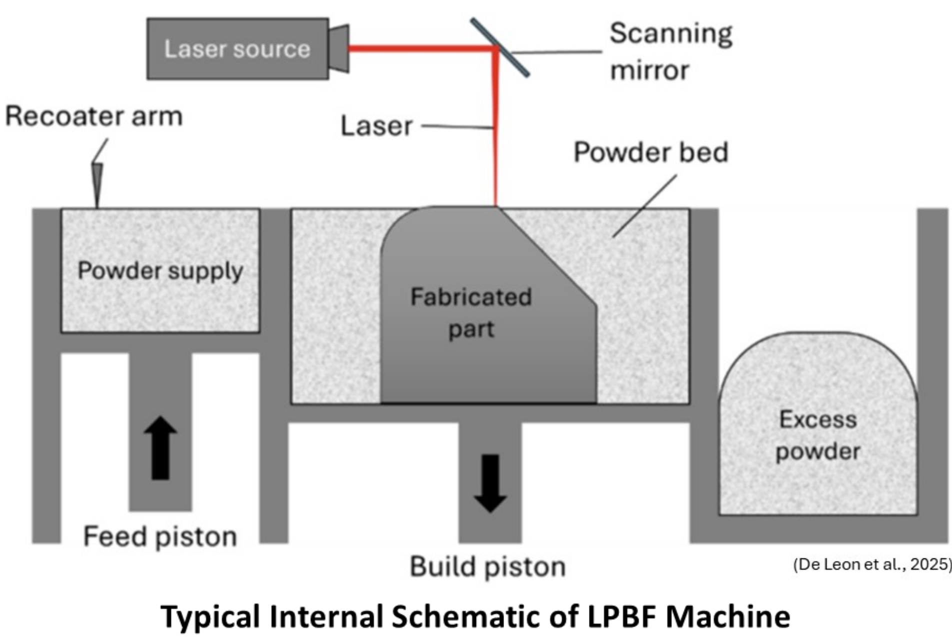

Figure 2.1. LPBF process schematic. A feed piston exposes powder to the surface where a recoater arm spreads the powder across the build plate. Excess powder is dumped into the storage container, and the laser source selectively melts the powder. Subsequent layers are spread and melt to build the part.

This chapter reviews the scientific and technical background required for localized XCT-labeled pore prediction in LPBF. The review is organized around the measurement-to-label-to-model pipeline used in this dissertation. First, pore formation mechanisms in LPBF are reviewed, with emphasis on lack-of-fusion porosity, gas porosity, and keyhole-related porosity. Second, post-process and in-situ methods for observing defects are reviewed. Third, the in-situ sensor modalities used in this work are discussed, including pyrometry, coaxial melt-pool imaging, acoustic emission, off-axis imaging, photometric-stereo heightmaps, and process data. Fourth, the need for synchronization, machine-coordinate mapping, and XCT-to-process registration is discussed. Fifth, prior machine-learning approaches for LPBF defect prediction are reviewed. Finally, multi-sensor fusion, physics-informed learning, and the research gap addressed by this dissertation are summarized.

The purpose of this chapter is not to review all LPBF monitoring studies. Instead, the goal is to identify the literature required to justify the specific approach used in this dissertation. The central argument is that localized pore prediction requires more than a sensor and a classifier. It requires synchronized sensor data, spatial registration to machine coordinates, three-dimensional pore ground truth, localized labels, and a fusion method capable of combining heterogeneous sensor inputs (Everton et al.; Grasso and Colosimo; Gorgannejad et al.; Satterlee et al.).

2.2 LPBF Process Physics and Defect Formation

LPBF produces metal parts through repeated melting and solidification of thin powder layers. The laser-material interaction is controlled by laser power, scan speed, spot size, hatch spacing, layer thickness, powder absorptivity, thermal properties, and local powder packing. These parameters influence the melt-pool temperature field, melt-pool dimensions, cooling rate, vaporization, recoil pressure, Marangoni convection, and solidification behavior. As a result, small changes in processing conditions can produce significant changes in melt-pool geometry and defect formation (King et al.; DebRoy et al.; Bidare et al.).

Porosity in LPBF can form through multiple mechanisms. Lack-of-fusion pores can form when insufficient energy input prevents complete melting or bonding between adjacent tracks or layers. These pores are commonly associated with low energy input, poor overlap, insufficient remelting, local powder defects, or irregular layer spreading. Gas pores can originate from entrapped gas in the powder, shielding-gas interactions, or gas entrapment during melting. Keyhole pores can form when high energy density produces a deep and narrow vapor depression in the melt pool. If the vapor depression becomes unstable, it can collapse and leave a bubble that becomes trapped during solidification (Huang et al.; Zhao et al.).

The distinction between these mechanisms is important, but the source of an individual pore cannot always be determined from XCT morphology alone. A pore detected in the final part may be influenced by keyhole instability, powder-bed variation, gas entrapment, scan-path boundary conditions, lack of fusion, or an interaction between several process variables. Therefore, the present dissertation treats XCT-detected porosity as the primary label and uses process signatures and sensor features to evaluate whether different pore populations can be predicted (King et al.; Khairallah et al.; Huang et al.; Zhao et al.).

2.3 Keyhole-Related Porosity and Melt-Pool Instability

Keyhole-related porosity is one of the main physical motivations for this dissertation. Under high-energy-density conditions, the laser can create a deep vapor depression in the melt pool. This depression increases absorptivity by allowing multiple internal reflections, but it can also become unstable. When the keyhole fluctuates, pinches, or collapses, vapor bubbles can be generated. If these bubbles are not removed before solidification, pores remain in the final part (Huang et al.; Zhao et al.; Khairallah et al.).

Operando synchrotron X-ray imaging has been especially important for identifying these mechanisms. Keyhole fluctuation, collapse, bubble formation, bubble growth, and bubble shrinkage have been directly observed during LPBF, and keyhole-pore formation has been connected to the dynamic behavior of vapor depression (Huang et al.; Martin et al.). A critical instability at the moving keyhole tip has also been shown to generate porosity during laser melting, with keyhole porosity associated with deep and narrow vapor depressions under high-power and low-scan-speed conditions as shown in Figure 2.2 (Zhao et al.; Cunningham et al.).

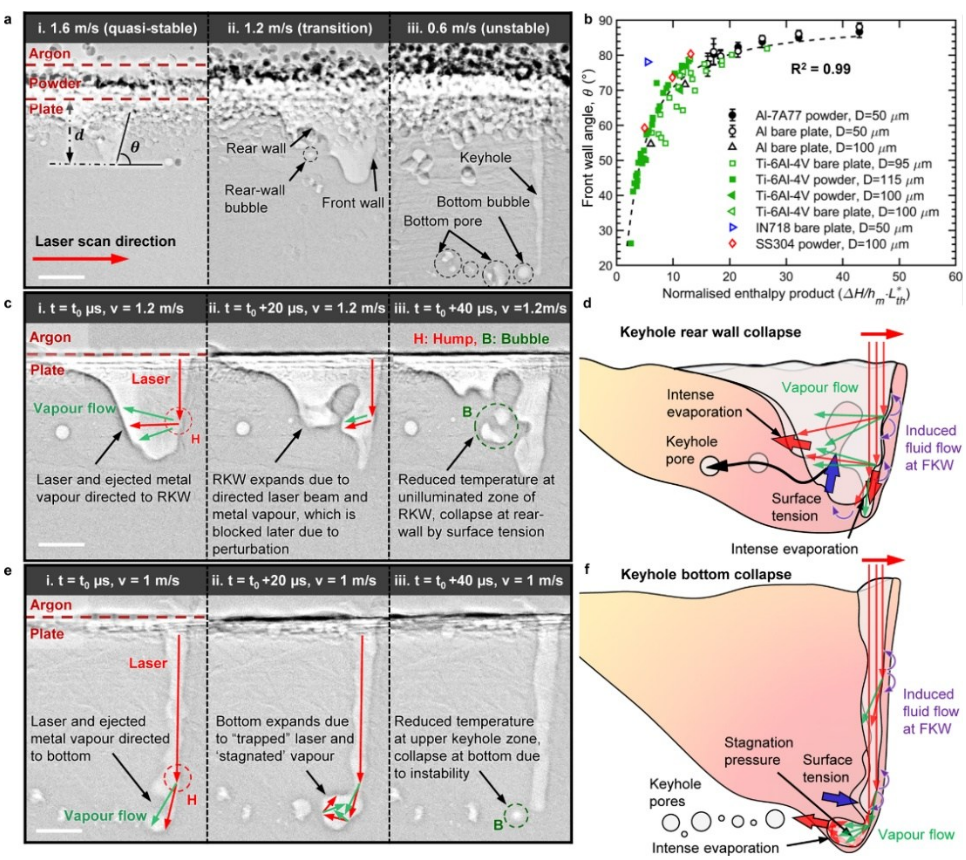

Figure 2.2. Keyhole collapse mechanism and keyhole melting regime transition. Keyhole morphology changes from wide and shallow to narrow and deep as the process moves from quasi-stable, transition, and unstable keyhole regimes. Pore formation can occur when keyhole instability or rear-wall collapse produces bubbles that are trapped by the solidification front (Huang et al).

These results are important because they show that pore formation can depend on transient melt-pool behavior rather than only on average process parameters. This supports the use of high-rate thermal, optical, and acoustic sensing. However, synchrotron imaging is not practical for routine manufacturing. Therefore, deployable sensors must be used to infer related process signatures during real LPBF builds (Huang et al.; Zhao et al.; Gorgannejad et al.; Martin et al.).

In this dissertation, keyhole-related physics is used as one important motivation for sensor selection and physics-informed feature design. Melt-pool geometry, pyrometer intensity, pyrometer derivatives, acoustic activity, and topographic features are all treated as indirect measurements of pore-relevant process behavior. However, it is not assumed that all pores in the dataset are keyhole pores.

2.4 Post-Process Characterization and Ground Truth

Post-process characterization is required to determine whether pores are present in the final part. Common methods include Archimedes density, metallographic cross-sectioning, mechanical testing, and X-ray computed tomography. Archimedes density provides a bulk density measurement but does not localize individual pores. Cross-section imaging can reveal pores and microstructure on polished surfaces, but it is destructive and only samples selected planes. Mechanical testing provides final property information, but it does not localize defect formation events (Everton et al.; Grasso and Colosimo; Oster et al.).

XCT is especially useful for pore-label generation because it provides a three-dimensional measurement of the internal pore field (Thompson et al.). Individual pores can be segmented, localized, and characterized by volume, equivalent diameter, centroid, bounding box, and shape metrics. However, XCT also introduces limitations. Resolution limits determine the smallest pores that can be detected. Segmentation thresholds can influence pore size and count. Beam-hardening, reconstruction artifacts, and surface effects can influence segmentation. Most importantly for this dissertation, the XCT coordinate system must be registered to the machine coordinate system before XCT pores can be used as labels for in-situ sensor data (Oster et al.; Gorgannejad et al.; Kim et al.).

Synchrotron X-ray imaging provides a different type of ground truth. Instead of measuring the final pore field after the build, synchrotron imaging can observe pore formation during laser melting. This enables direct temporal registration between pore-formation events and sensor signals. Multimodal process monitoring with synchrotron X-ray radiography has been used to predict localized keyhole pore formation over millisecond-scale windows (Gorgannejad et al.). However, synchrotron imaging is not deployable for routine manufacturing. It is best interpreted as a mechanism-discovery and validation tool. XCT, in contrast, is a practical ex-situ ground-truth method for training and validating deployable in-situ monitoring systems.

2.5 In-Situ Sensors for LPBF Pore Monitoring

In-situ monitoring is used to measure process signatures during printing. The goal is not to observe pores directly, but to measure process behavior that may correlate with pore formation. Multiple sensor types have been investigated in LPBF, including thermal cameras, pyrometers, photodiodes, coaxial cameras, off-axis cameras, acoustic emission sensors, layer-wise optical imaging, profilometry, and process-log data. Each sensor provides different information and has different limitations (Everton et al.; Grasso and Colosimo; Tapia and Elwany).

2.5.1 Pyrometry and Thermal Emission

Pyrometers measure optical or infrared emission from the melt-pool region. Thermal emission is relevant because pore formation depends strongly on local energy input, temperature history, vaporization, and cooling behavior (Hooper). In LPBF, thermal monitoring can be used to infer process instability, local heat accumulation, lack-of-fusion risk, or keyhole-related behavior (Diehl et al.; Everton et al.; Grasso and Colosimo).

However, pyrometer signals must be interpreted carefully. The measured signal is affected by emissivity, wavelength response, viewing geometry, melt-pool size, plume emission, sensor bandwidth, and calibration. Without thermal calibration, pyrometer channels should not be treated as direct absolute temperature measurements. Instead, relative intensity, channel ratios, derivatives, and normalized features can be used as physically motivated indicators of changes in thermal behavior. Ratiometric and bichromatic pyrometry have been used to relate photodetector signals to temperature, but calibration and emissivity remain major considerations (Diehl et al.; Lane et al.).

In this dissertation, three pyrometer channels are used as high-rate thermal sensors. The channels are treated primarily as relative thermal-emission signals. This supports the use of derivative-based and ratio-based features, especially when physics-informed features such as relative temperature proxies or melt-pool-depth derivatives are evaluated (Diehl et al.).

2.5.2 Coaxial Melt-Pool Imaging

Coaxial melt-pool imaging observes the melt region through or near the laser optical axis. This sensor modality is important because pore formation is directly related to melt-pool geometry and stability. Melt-pool area, intensity distribution, aspect ratio, major/minor axis length, and frame-to-frame changes can provide information about conduction-mode melting, keyhole-mode melting, spatter, plume behavior, and unstable process conditions (Huang et al.; Zhao et al.; Gorgannejad et al.; Clijsters et al.).

Coaxial imaging has a stronger spatial relationship to the melt event than many off-axis or global sensors because the field of view follows the laser interaction region. In this dissertation, the coaxial camera is included to measure melt-pool geometry and image-based process signatures. The coaxial camera was included because it provides a direct optical view of the laser–material interaction region.

The limitation of coaxial imaging is that image intensity and apparent melt-pool shape can be affected by exposure, saturation, plume emission, optical alignment, frame rate, thresholding, and the optical path. Therefore, the images are used both as raw image inputs and as a source of geometric features, rather than as direct measurements of melt-pool depth or pore formation (Everton et al.; Grasso and Colosimo).

2.5.3 Acoustic Emission

Acoustic emission (AE) sensors measure high-frequency elastic waves produced by rapid energy release or dynamic material response. In LPBF, AE may be influenced by vaporization, spatter, cracking, keyhole collapse, melt-pool instability, and mechanical response of the build plate or part. AE is attractive because it can capture dynamic events that may not be visible in thermal or optical images (Gorgannejad et al.; Everton et al.; Grasso and Colosimo).

AE has been used in prior work for detecting and classifying defects or process states in laser-based manufacturing, and multimodal process monitoring with AE has been used for keyhole-pore prediction under synchrotron observation (Gorgannejad et al.; Shevchik et al.). In the present dissertation, AE is used as a high-rate one-dimensional signal synchronized to laser position. AE was included because it may capture transient process behavior not represented by thermal or optical signals.

The limitation of AE is that the measured signal is affected by sensor coupling, propagation path, machine noise, chamber acoustics, recoater motion, and build-plate response. As a result, AE is not interpreted as a direct pore measurement. It is treated as a dynamic process signature that may become more useful when fused with thermal, optical, and process information (Gorgannejad et al.; Everton et al.).

2.5.4 Off-Axis Imaging and Photometric-Stereo Heightmaps

Off-axis imaging provides layer-wise information about the powder bed and top surface. This information is important because defects can be influenced by the condition of the layer before melting. Streaks, humps, depressions, recoater marks, spatter deposits, and uneven spreading can alter the local powder layer and change the effective energy input. These features may act as defect precursors even if they are not visible in the melt-pool camera (Everton et al.; Grasso and Colosimo; Woodham; Scime and Beuth).

Photometric stereo is used to estimate surface orientation from images captured under different illumination directions while keeping the viewing direction fixed (Woodham). In LPBF, this concept can be adapted to multi-light powder-bed imaging. Directional illumination is used to capture multiple off-axis images, surface normals are estimated, and a heightmap is generated. The heightmap can then be mapped to machine coordinates and used as a topographic input to the model.

In this dissertation, raw off-axis images and photometric-stereo heightmaps are used as part of the multimodal input. These data provide layer-surface and powder-bed information that is complementary to melt-pool imaging, pyrometry, AE, and process data. The limitation is that photometric-stereo reconstruction depends on lighting calibration, camera calibration, surface reflectance assumptions, shadowing, saturation, and image-to-machine coordinate mapping. These limitations are addressed through calibration and by treating the heightmap as a relative topographic feature rather than an absolute density measurement unless validated separately (Woodham; Everton et al.).

2.6 Process Parameter Data

Process parameter data provides the spatial and operational context required for localized pore prediction. Laser position, laser drive, layer number, scan speed, laser power, hatch spacing, spot size, and layer height are required to determine where and under what nominal condition a sensor signature was collected. Process data also provides the reference frame for synchronizing one-dimensional signals, image data, and XCT labels (Yeung and Grantham; Hocker et al.).

Process parameters are not sufficient to predict all pores because local powder condition, thermal history, edge effects, and stochastic melt-pool behavior also influence defect formation. However, process parameters remain important because they define the operating regime. A sensor signature may have different meanings under different laser power or scan speed conditions. For this reason, process parameters are used both for synchronization and as inputs for process-informed PMMSF (King et al.; DebRoy et al.; Khairallah et al.).

2.7 Machine Learning-based LPBF Defect Prediction

Machine learning has been increasingly used for LPBF process monitoring because the relationship between in-situ sensor signals and final defects is high-dimensional and nonlinear. Prior work has used thermal images, optical tomography, photodiode signals, acoustic signals, process data, and fused sensor features to detect defects, classify process regimes, predict porosity, or estimate part quality (Everton et al.; Grasso and Colosimo; Tapia and Elwany; Oster et al.; Gorgannejad et al.).

Thermal monitoring is one of the most common approaches (Baumgartl et al.). Short-wave infrared thermography has been used to reconstruct thermal history and predict porosity-related defects in LPBF. In one deep-learning framework, thermographic data were used to predict local porosity with XCT-based labels, demonstrating that in-situ thermal history can contain defect-relevant information (Oster et al.). This work achieved an F1 score of 86% for pore prediction but required a relatively large voxel size of 700μm x 700μm x 50μm. Other work has used optical or thermal signatures with machine learning to detect pores or anomalies in LPBF (Galarza et al.). This work produced an accuracy of 85% but required a 1.5ms signature window, corresponding to approximately 1 mm at the applicable laser speed.

Multimodal approaches have also been investigated. Localized keyhole-pore prediction has been demonstrated using multimodal process monitoring and X-ray radiography, where thermal and acoustic signals were temporally registered to pore formation events observed by high-speed synchrotron X-ray imaging as shown in Figure 2.3 (Gorgannejad et al.). This work produced an accuracy of 90% on pore prediction but was applied to single scan lines and thus does not account for multi-layer or neighboring scan effects.

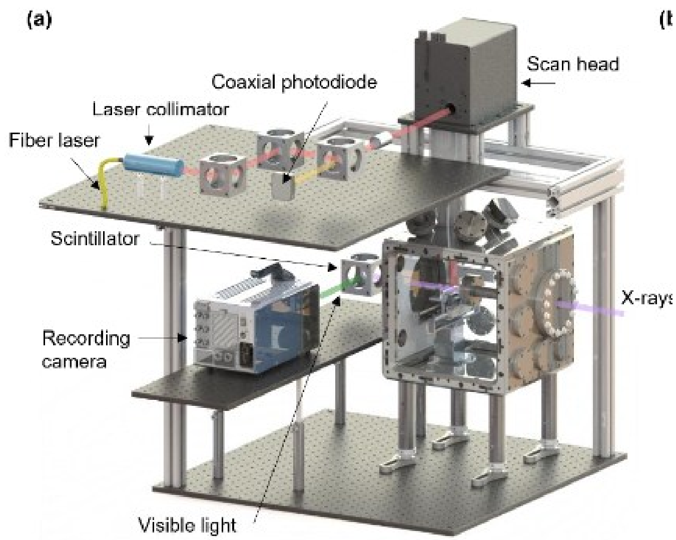

Figure 2.3. Registering in-situ process signals with synchrotron X-ray imaging as proposed by Gorgannejad et al

More recent work has extended in-situ sensor monitoring toward multi-class gas-porosity formation, highlighting the need to move beyond binary keyhole-pore classification (Giri et al.). Their framework classified no-pore, small-pore, and large-pore conditions and achieved 79% accuracy and an ROC-AUC score of 0.89 with five-fold cross-validation. Infrared and dual-scan-head methods shown in Figure 2.4 have also been developed to observe cooling behavior and detect controlled (implanted) defects, although defect size, artificial defect geometry, or measurement configuration may limit direct transfer to stochastic pore prediction in production LPBF (Höfflin et al.). Their study resulted in an accuracy of 93% on 700µm defects but reduced to an accuracy of 67% on 500µm defects.

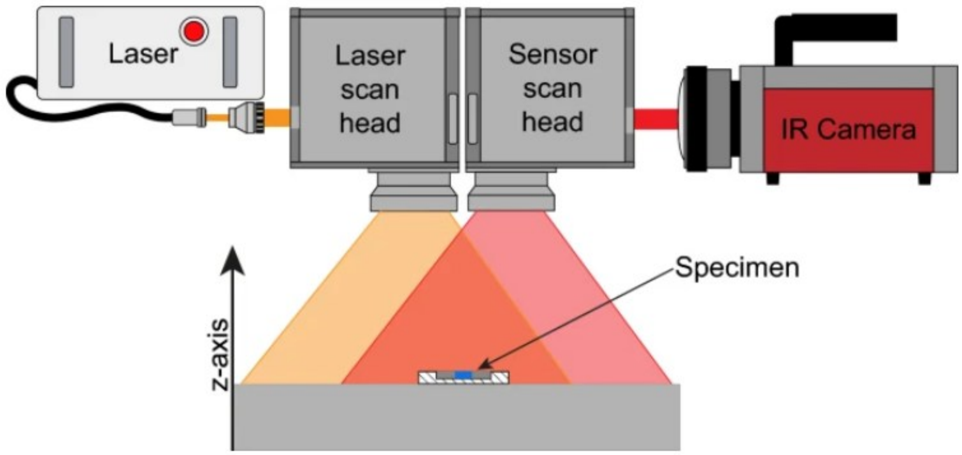

Figure 2.4. Dual head setup proposed by Höfflin et al for in situ defect detection (Höfflin et al.)

These studies demonstrate that in-situ sensor data can be related to pore formation. However, several limitations remain. First, some studies rely on synchrotron or X-ray radiography during laser melting, which provides excellent ground truth but is not deployable for routine production monitoring. Second, some studies use implanted or artificial voids, which may not represent naturally occurring stochastic pores. Third, large analysis windows may increase performance while reducing spatial precision. Fourth, many studies use one primary sensor modality or a limited fusion strategy. Fifth, localized XCT-labeled pore prediction requires spatial registration between final pore locations and process-centered sensor windows, which is not always the focus of prior work (Oster et al.; Gorgannejad et al.; Giri et al.; Höfflin et al.).

The present dissertation is positioned within this gap. XCT is used as ex-situ ground truth, but deployable sensors are used for prediction. Labels are generated locally by registering the XCT pore field to machine coordinates. Multiple sensor modalities are used because pore-relevant process signatures may appear in thermal, optical, acoustic, topographic, and process-parameter data.

2.8 Sensor Synchronization and Coordinate Registration

Localized pore prediction requires that each sensor input and each pore label refer to the same physical region. This requirement makes synchronization and coordinate registration central to the problem. One-dimensional sensors such as AE and pyrometers are recorded as time series. Coaxial images are recorded as high-speed frames. Off-axis images and heightmaps are recorded in pixel coordinates. XCT pores are measured in voxel coordinates. Process data are recorded as scan rows with laser position and layer number. These data streams must be transformed into a common coordinate framework before supervised learning can be performed (Hocker et al.; Yeung and Grantham; Kim et al.).

Laser-trigger signatures provide a practical synchronization reference. If the laser trigger or laser drive is captured in the process data and in the sensor signals, then sensor windows can be mapped to laser-on process rows. Laser position provides the corresponding X/Y location, while layer number provides the Z location. Off-axis images require an additional calibration step because image pixels must be mapped to laser or machine coordinates. XCT requires a separate registration step because the reconstructed volume must be converted from voxel coordinates to millimeters and aligned to the machine coordinate system (Yeung and Grantham; Hocker et al.; Kim et al.).

This step is often described as preprocessing, but it is a major source of uncertainty in localized pore prediction. If the registered XCT pore field is shifted relative to the laser path, then labels are assigned to the wrong sensor windows. This can create false positives, false negatives, and misleading model performance (Hocker et al.; Yeung and Grantham; Kim et al.).

2.9 Multi-Sensor and Multimodal Fusion

Multi-sensor fusion is motivated by the fact that different sensors measure different physical phenomena. In LPBF, pore formation may be associated with thermal history, melt-pool geometry, acoustic transients, powder-bed condition, local scan path, and process parameters. A single sensor may detect only part of this behavior. Therefore, a fusion method is needed to combine complementary sensor information (Baltrusaitis et al.; Satterlee et al.; Everton et al.).

Simple fusion methods include input concatenation, channel fusion, feature concatenation, and head fusion. Input concatenation combines sensor data before or near the beginning of the model. This approach can be difficult when sensors have different sampling rates, dimensions, and noise characteristics. Feature concatenation combines extracted features from different sensors, but the fusion may occur only after sensor-specific interactions have already been compressed. Head fusion trains separate branches and combines final predictions or final feature vectors. This preserves sensor-specific feature extraction but may miss intermediate interactions between sensors (Baltrusaitis et al.; Satterlee et al.).

PMMSF is developed to address limitations of simple fusion approaches. PMMSF processes signals through sensor-specific Feature Extraction Networks (FEN) and repeatedly fuses through a Sensor Fusion Network (SFN). This architecture allowed heterogeneous sensor streams to contribute at multiple network depths rather than only through an initial concatenated input or a final fused embedding. PMMSF was originally developed for multimodal LPBF monitoring, while its first published implementation and experimental validation were performed for multi-sensor pipe-leak detection. The pipe-leak results demonstrated the viability of the progressive fusion architecture before its application to the LPBF pore-prediction problem investigated in this dissertation. The method achieved 97.1% leak-versus-non-leak detection accuracy and leak-location classification accuracies between 95.5% and 97.4% for fused sensor configurations, providing the methodological basis for adapting the architecture to LPBF pore prediction (Satterlee et al.).

The term multimodal is used because the LPBF inputs include one-dimensional signals, image data, heightmap/topographic data, process parameters, and physics-informed features. This adaptation is necessary because LPBF pore prediction requires fusion of sensors that differ not only in sampling rate, but also in data structure and physical meaning (Baltrusaitis et al.; Satterlee et al.).

PMMSF is expected to be useful for LPBF because weak sensors may still provide useful information when combined with stronger sensors. For example, pyrometer channels may not be the strongest individual predictors, but they may provide thermal information that improves interpretation of coaxial image or AE features. Similarly, off-axis heightmaps may not directly identify pores, but they may provide layer-condition information that changes pore risk in later layers. Layer-wise fusion allows these contributions to be learned across the feature hierarchy rather than requiring all useful information to appear in the final feature vector (Baltrusaitis et al.; Satterlee et al.).

2.10 Physics-Informed and Process-Informed Learning

Physics-informed machine learning is used when physical knowledge is incorporated into a learning model. This can be done through physics-based losses, auxiliary outputs, feature engineering, attention or gating, graph structures, or constraints on model predictions. In LPBF, physical information is useful because the available datasets are often limited and the process is governed by known thermal, fluid, and geometric relationships (Karniadakis et al.; King et al.; DebRoy et al.).

In this dissertation, physics-informed learning is focused on feature-informed and fusion-level-informed approaches. Physics-derived features are used as inputs or fusion-guidance variables rather than as hard constraints in the loss function. These features include melt-pool-depth proxies, melt-pool-depth derivatives, relative temperature proxies, melt-pool geometry, and topographic features. The objective is not to replace the sensor data with a simplified physics model. Instead, physics-derived features are used to help the model interpret the sensor data in a more process-relevant way (Karniadakis et al.; Huang et al.; Zhao et al.; Diehl et al.).

Melt-pool depth is one physics-informed quantity of interest because deeper and unstable melt pools are associated with keyhole behavior. Direct melt-pool-depth measurement is not available from the production sensor suite, so empirical or proxy estimates must be used. A melt-pool-depth derivative is especially useful because pore formation is expected to depend on instability and dynamic changes, not only on the absolute value of the estimated depth. This is also practical because uncalibrated pyrometers may be more reliable as relative or derivative signals than as absolute temperature measurements (Huang et al.; Zhao et al.; Diehl et al.).

Relative temperature proxies are also considered because multiple pyrometer channels can provide information about thermal emission at different wavelengths. Bichromatic or radiometric thermometry can relate signal ratios to temperature under specific assumptions, but calibration and emissivity effects must be considered (Diehl et al.). Therefore, relative temperature is treated as a proxy rather than as a fully calibrated absolute temperature.

Process-informed learning is also used because the meaning of a sensor signature can depend on the operating condition. Laser power, scan speed, hatch spacing, layer height, spot size, volumetric energy density, and areal energy density define the nominal process regime. A sensor signal that indicates high pore risk under one condition may have a different meaning under another. In process-informed PMMSF, process parameters can be concatenated with features, passed through a weighting network, or used to guide SoftMax-based fusion weights (King et al.; DebRoy et al.; Karniadakis et al.).

Direct physics-based modeling of keyhole-collapse frequency using Rayleigh-Plesset-type equations is not implemented as a main feature in this dissertation. It remains a future direction because the assumptions required to connect acoustic emission signals to collapsing vapor cavities are not yet validated for the current sensor configuration and dataset.

2.11 Research Gap

The literature supports several conclusions. First, LPBF part quality is limited by process-induced defects and by the difficulty of connecting process parameters to final structure and properties. Second, keyhole fluctuation and melt-pool instability are important pore-formation mechanisms, but not all pores can be assigned to a single mechanism. Third, in-situ monitoring can provide defect-relevant information, but single-sensor approaches may not capture all relevant physical phenomena. Fourth, X-ray synchrotron radiography provides strong mechanistic ground truth, but it is not practical for routine manufacturing deployment. Fifth, XCT provides practical three-dimensional ex-situ ground truth, but it must be registered to process coordinates before localized labels can be generated. Sixth, simple fusion methods may not be sufficient for heterogeneous LPBF sensors because the sensor streams differ in sampling rate, dimensionality, and physical meaning (King et al.; DebRoy et al.; Huang et al.; Zhao et al.; Everton et al.; Grasso and Colosimo; Oster et al.; Gorgannejad et al.; Satterlee et al.).

To fill this gap, a localized, XCT-labeled, multimodal sensor-fusion framework for LPBF pore prediction is proposed in this study. This framework must connect in-situ sensor signals to machine coordinates, register XCT pore morphology to the process coordinate system, generate local pore/no-pore labels, and fuse heterogeneous sensor information in a way that preserves weak but physically meaningful sensor contributions.

2.12 Chapter Summary

This chapter reviewed the background needed to support localized XCT-labeled pore prediction in LPBF. Porosity was shown to arise from multiple mechanisms, including lack of fusion, gas entrapment, keyhole instability, and powder-bed variation. Keyhole-related porosity was reviewed as a major motivation because melt-pool depth, vapor depression instability, and bubble entrapment have been directly connected to pore formation. Post-process and in-situ characterization methods were compared, with XCT identified as a practical ground-truth method and synchrotron imaging identified as a mechanism-discovery tool (King et al.; DebRoy et al.; Huang et al.; Zhao et al.; Khairallah et al.; Gorgannejad et al.).

The sensor modalities used in this dissertation were then reviewed. Pyrometry provides relative thermal emission, coaxial imaging provides melt-pool geometry, AE provides dynamic process signatures, off-axis imaging and photometric-stereo heightmaps provide layer-surface information, and process data provide spatial and operational context. Because these data are heterogeneous, synchronization and coordinate registration are required before localized learning can be performed (Everton et al.; Grasso and Colosimo; Diehl et al.; Woodham; Hocker et al.).

Prior machine-learning studies demonstrate that in-situ data can predict pores or defects, but existing approaches often rely on single sensors, large windows, artificial defects, synchrotron-based labels, or limited fusion methods. The research gap addressed by this dissertation is the development of a practical, localized, XCT-labeled, multimodal sensor-fusion framework for LPBF pore prediction. The next chapter defines the feature-engineering and sensor-fusion approach used to address this gap (Oster et al.; Gorgannejad et al.; Giri et al.; Satterlee et al.).

Chapter 3 Approach: Feature Engineering and Sensor Fusion

3.1 Overview

This chapter defines the modeling approach used to convert synchronized multimodal LPBF data and registered XCT pore labels into a supervised learning problem. The chapter connects the data-generation workflow to the model-evaluation workflow by specifying the localized prediction target, the engineered feature groups, the multi-layer context, the PMMSF architecture, and the comparison methods used to evaluate whether process-informed or physics-informed fusion improves pore prediction.

The approach is designed around a central limitation in localized LPBF monitoring: in-situ sensor streams contain indirect signatures of pore-relevant process behavior, but those signatures are heterogeneous in sampling rate, dimensionality, coordinate basis, and physical meaning. Feature engineering and sensor fusion are therefore treated as part of the research contribution rather than as incidental preprocessing. The feature-engineering workflow converts registered process, sensor, image, topographic, and pore-label data into model inputs. The sensor-fusion workflow then evaluates how those inputs should be combined to predict localized XCT-detected pore labels (Baltrusaitis et al.; Satterlee et al.; Karniadakis et al.).

The chapter is limited to the learning approach. Experimental hardware is described in Chapter 4. Coordinate mapping, sensor synchronization, XCT pore segmentation, XCT-to-process registration, and query-volume label generation are described in Chapter 5. Results from the trained models are reported in Chapter 6.

3.2 Localized Prediction Target

The supervised learning target is localized prediction of XCT-detected porosity. Each candidate process location is associated with a process-centered query volume after XCT pore morphology has been registered into the process-coordinate frame. A positive learning label is assigned when the registered center of an accepted XCT pore falls within the process-centered query volume; otherwise, the sample is assigned to a negative label. This definition differs from part-level porosity classification because the model is trained to associate sensor signatures with spatially localized pore regions rather than with an entire coupon or process-parameter group.

3.3 Sensor-Fusion Strategy

Feature fusion is used to make heterogeneous sensor streams comparable at a localized process location. The available inputs include one-dimensional time-series signals, coaxial melt-pool images, off-axis images, photometric-stereo heightmaps, process-coordinate fields, nominal process parameters, and registered pore-label metadata. The raw sensor data is fused by first localizing the signals in the same time domain or spatial domain and then selecting contiguous signals that fit within the desired window.

The one-dimensional signal inputs include the High, Low, and TED thermal-emission channels, acoustic emission, laser drive or trigger information, and derived local signal windows. The image inputs include raw or cropped coaxial melt-pool frames associated with the process location. The off-axis inputs include registered image products and photometric-stereo heightmaps that provide surface and powder-bed context. Process inputs include laser power, scan speed, layer height, hatch spacing, nominal laser spot diameter, volumetric energy density, areal energy density, part identifier, layer number, and scan-location metadata when available.

The feature-engineering workflow preserves the distinction between raw sensor inputs and derived physics or process features. Raw sensor inputs allow the neural network to learn data-driven representations, while engineered features provide compact process-relevant quantities that may improve robustness, interpretability, or fusion weighting.

3.4 Parallel Multi-Layer Multimodal Sensor Fusion

PMMSF was designed to overcome limitations of simple channel fusion, input concatenation, and final-layer head fusion by allowing sensor-specific feature extraction and repeated feature interaction across network depth as shown in Figure 3.1. The input branches may contain one-dimensional process signatures, coaxial images, off-axis images, photometric-stereo heightmaps, process parameters, and physics-informed features. Each input is processed by a branch appropriate for its dimensionality, and branch outputs are projected or pooled to compatible feature dimensions so that fusion can occur repeatedly during feature extraction rather than only at the classifier head.

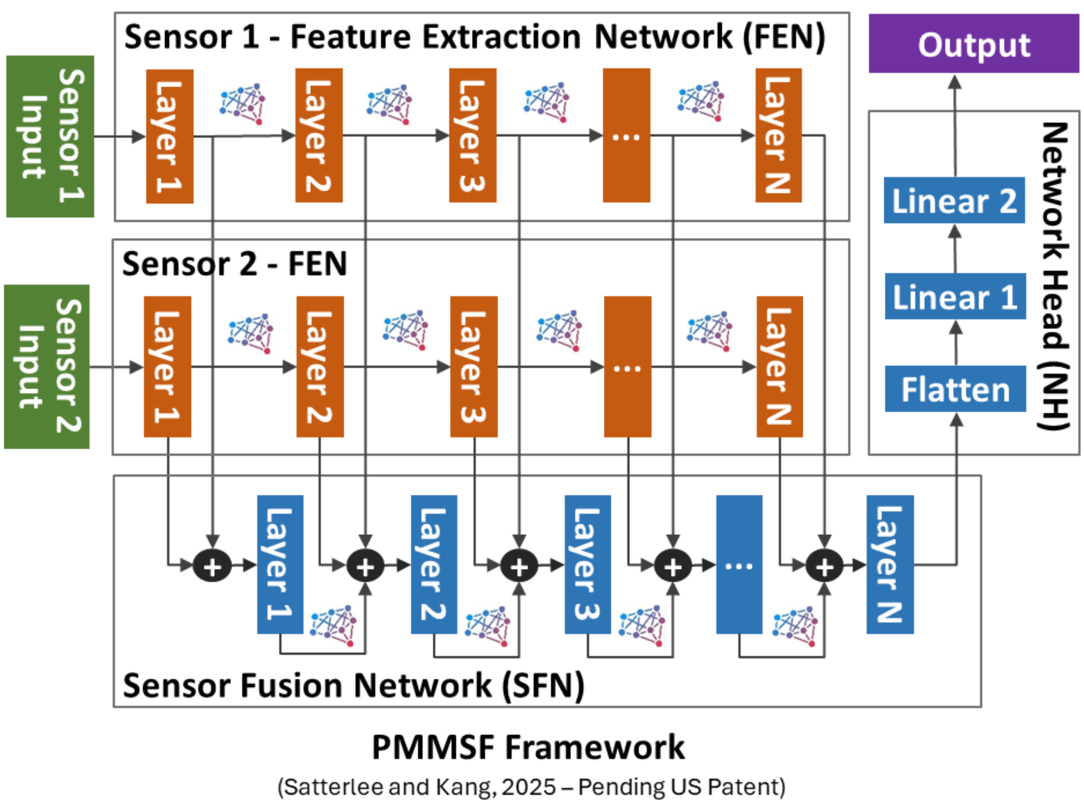

Figure 3.1. Architecture for PMMSF model. Each FEN network extracts features from the sensors, and features are combined within the SFN.

The PMMSF formulation begins with a sensor-specific Feature Extraction Network (FEN). For modality or sensor j at layer i, the one-dimensional convolutional branch update is

$$y_{i,j}(m) = \sum_{n=0}^{k-1} y_{i-1,j}(m+n) w_{i,j}(n) \tag{3.1}$$

where y is the layer-specific representation for sensor j, m is the convolution-output index, k is the convolution-kernel length, and w is the corresponding sensor-branch convolutional kernel. Equation 3.1 states that each branch first learns modality-specific features before those features are merged with features from the other branches.

For the two-sensor case presented in the PMMSF publication, the first fusion layer combines the first-layer sensor features and applies a fusion-network kernel:

$$f_1(m) = \sum_{n=0}^{k-1} [y_{1,1}(m+n) + y_{1,2}(m+n)] w_{f,1}(n) \tag{3.2}$$

From the second fusion layer onward, the Sensor Fusion Network (SFN) combines current sensor-branch features with the previous fusion representation. Generalized to s LPBF modalities, the layer-wise fusion update is

$$f_1(m) = \sum_{p=0}^{k-1} [\sum_{j=1}^{s} y_{i,j}(m+p) + f_{i-1}(m+p)] w_{f,1}(p) \tag{3.3}$$

For two sensors, expansion of the recursive fusion term shows how the current fusion state depends on both the sensor-branch kernels and the previous fusion kernel

$$f_1(m) = \sum_{p=0}^{k-1} \left[\sum_{n=0}^{k-1} y_{i-1,1}(m+p+n) \left(w_{i-1,1}(n) + w_{f,i-1}(n)\right) + \sum_{n=0}^{k-1} y_{i-1,2}(m+p+n) \left(w_{i-1,2}(n) + w_{f,i-1}(n)\right)\right] w_{f,1}(p) \tag{3.4}$$

This expanded form explains why PMMSF differs from final-head fusion. Sensor information is not independently compressed and then merged only once. Instead, the fusion state is updated across the feature hierarchy, allowing the model to promote, demote, or combine sensor-specific features before the final embedding is formed (Satterlee et al.).

In the LPBF implementation, the FEN branch operator in Equation 3.1 is interpreted broadly. For one-dimensional signals, the operator is a one-dimensional convolutional feature extractor. For images or heightmaps, 2D convolutional operators are used prior to linearizing features. For scalar or engineered vectors, it is a feature projection. The common requirement is that each branch produces an intermediate representation that can be supplied to the SFN at a compatible layer or stage.

Process-informed and physics-informed PMMSF extends the same layer-wise structure by allowing auxiliary variables to guide fusion weights. Let g denote process or process-informed conditioning variables, such as laser power, scan speed, layer height, or spot size. A gated PMMSF stage used in this dissertation can be represented as

$$\widetilde{f_i} = \varphi_i\left(\sum_{j=1}^{s} \alpha_{i,j}(g) P_{i,j}\left(h_{i,j}\right) + \alpha_{i,f}(g) P_{i,f}\left(f_{i-1}\right)\right) \tag{3.5}$$

where h is the layer-specific branch representation, the P terms are projection or pooling operators that place branch features and the prior fusion state into compatible representations, the alpha terms are learned conditioning weights, g is the gate features, and Phi is the fusion-stage transformation. When process or physics conditioning is not used, the alpha terms are reduced to learned fusion weights that depend only on trainable model parameters. When conditioning is enabled, the model can learn that a branch should be weighed differently under different process regimes or local physical states.

The PMMSF training schedule is staged so that the feature extractors and fusion network do not have to learn all structure simultaneously. First, each enabled feature-extraction branch is pretrained independently so that each modality develops a useful representation before fusion.

After pretraining, the feature-extraction networks and the fusion network are trained alternately with separate Adam optimizers: during feature-network epochs, the fusion parameters are held fixed so that the modality branches align their representations with the fusion pathway; during fusion-network epochs, the feature branches are held fixed so that the SFN learns how to weight and combine the available branch features. This alternating schedule follows the original PMMSF rationale that the FEN learns sensor-specific features while the SFN learns how to combine those features across layers (Satterlee et al.).

After the alternating stage, the training schedule uses three refinement phases. The feature-extraction branches are fine-tuned while the fusion network remains fixed, then the fusion network is fine-tuned while the feature-extraction branches remain fixed, and finally the full PMMSF model is fine-tuned jointly with a shared Adam optimizer. This sequence preserves the benefit of branch pretraining, reduces the risk that the fusion network learns unstable combinations before each modality has a meaningful representation, and then allows the final model to adjust the branches and fusion pathway together for the localized pore/no-pore objective.

This structure is used in the dissertation because localized LPBF pore prediction requires simultaneous preservation of modality-specific information and repeated cross-modality interaction. Input concatenation can force incompatible signals into a single representation too early, while head fusion can delay sensor interaction until each branch has already compressed its information. PMMSF preserves independent feature extraction while allowing layer-wise fusion to combine thermal, acoustic, optical, topographic, process-parameter, and physics-informed information before the final pore/no-pore prediction is made.

3.5 Physics-Informed Features

Physics-informed features are included because pore formation is influenced by local energy input, melt-pool geometry, melt-pool stability, thermal history, and surface condition. These features are not treated as exact physical measurements unless calibration supports that interpretation. Instead, they are used as physically motivated proxies that may help the model interpret sensor signatures under different process regimes (King et al.; DebRoy et al.; Huang et al.; Zhao et al.; Karniadakis et al.).

Pyrometer-derived features include relative temperature. Because the pyrometer channels are not treated as calibrated absolute-temperature measurements, derivative-based and ratio-based quantities are preferred when the goal is to capture relative thermal behavior. The temperature proxy is computed as the inverse-log of the ratios of the high (S1) and low sensor (S2) as shown below:

$$T = \frac{A}{\ln\left(\frac{B}{\frac{S_1(T)}{S_2(T)}}\right)} \tag{3.6}$$

Where A and B are calibrated constants. In this study pyrometers are uncalibrated, so they are set to 1 and only relative temperature is used.

Coaxial image features can include melt-pool area, aspect ratio, x and y position, contour, intensity statistics, shape descriptors, and frame-to-frame changes as shown in Figure 3.2. Acoustic features can include windowed amplitude, energy, spectral, or transient descriptors extracted over laser-on intervals. Off-axis and photometric-stereo features provide surface and topographic context from the current or previous layer (Diehl et al.; Gorgannejad et al.; Woodham; Everton et al.).

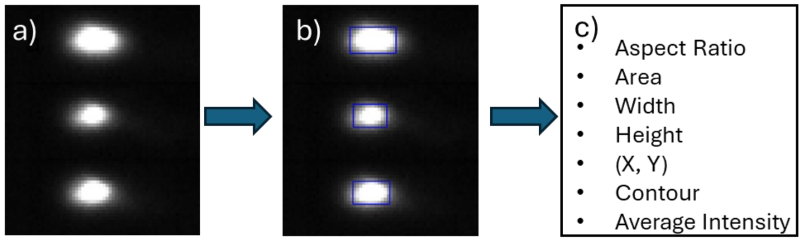

Figure 3.2. Contiguous raw coaxial camera images (a) are segmented using a threshold of 150 (b) to obtain melt-pool features (c). The sequence can then be used to compute frame-to-frame changes.

Melt-pool-depth proxy features are included as an exploratory physics-informed feature group. Direct melt-pool depth is not measured by the production sensor suite; therefore, melt-pool-depth estimates are treated as empirical, or proxy quantities derived from available thermal and image signatures. The derivative of a melt-pool-depth proxy is especially relevant because pore formation can be associated with process instability or rapid state changes rather than with a single static intensity value (Huang et al.; Zhao et al.; Diehl et al.; Hofmann et al). In Hofmann et al. a regression model was performed to determine the melt-pool depth as a function of laser power, P, scan speed, v, laser spot size, D, and powder layer thickness, t.

$$M = \exp(y) \tag{3.7}$$

$$y = c_1 + c_2P - c_3v - c_4D + c_5t + c_6P^2 + c_7v^2 + c_8D^2 + c_9t^2 + c_{10}Pv + c_{11}PD + c_{12}vD$$

$$dM = \frac{\partial M}{\partial P}dP + \frac{\partial M}{\partial v}dv + \frac{\partial M}{\partial D}dD + \frac{\partial M}{\partial t}dt \tag{3.8}$$

If layer thickness and velocity are assumed to be constant, then we are only left with power change and laser spot size change, and we can approximate it as:

$$dM \approx \frac{\partial M}{\partial P}dP + \frac{\partial M}{\partial D}dD \tag{3.9}$$

The laser spot size is obtained from the co-axial camera using the minor axis as shown in Figure 3.2. The derivative with respect to time is estimated using finite-difference.

$$\frac{dD}{dt} = \frac{D_2 - D_1}{\Delta t} \tag{3.10}$$

We assume that laser power is proportional to absorbed laser power. The power can be estimated from the Stefan–Boltzmann relationship by

$$P = \varepsilon \sigma A T^4 \tag{3.11}$$

Taking the derivative with respect to time we find

$$\frac{dP}{dt} = 4\varepsilon\sigma A T^3 \frac{dT}{dt} \tag{3.12}$$

Plugging Equation 3.11 back in relative rate becomes

$$\frac{1}{P}\frac{dP}{dt} = 4\frac{1}{T}\frac{dT}{dt} \tag{3.13}$$

The fractional rate of change in radiated power is four times the fractional rate of change in absolute temperature under the stated assumptions.. We can then compute the relative differential and insert the terms

$$\frac{dM}{M} \approx \frac{P}{M}\frac{\partial M}{\partial P}\frac{dP}{P} + \frac{D}{M}\frac{\partial M}{\partial D}\frac{dD}{D} \tag{3.14}$$

For a given operating point, we can use the process parameter inputs to estimate and linearize for that process parameter location:

$$S_p = \frac{P}{M}\frac{\partial M}{\partial P} \quad \text{and} \quad S_D = \frac{D}{M}\frac{\partial M}{\partial D} \tag{3.15}$$

$$\frac{dM}{M} \approx S_p \frac{dP}{P} + S_D \frac{dD}{D} \approx 4S_p \frac{dT}{T} + S_D \frac{dD}{D} \tag{3.16}$$

For a laser power of 180 W, nominal spot diameter of 80 μm, scan speed of 1050 mm/s, and layer height of 30 μm, the estimated melt-pool depth is $M \approx 45\ \mu$m. Substitution into Equation 3.16 gives the normalized sensitivity coefficients used to construct the melt-pool-depth derivative proxy.

$$\frac{dM}{M} \approx 4(0.965)\frac{dT}{T} - 1.08\frac{dD}{D} \approx 4\frac{dT}{T} - \frac{dD}{D}$$

Therefore, we can estimate the relative change in melt-pool depth using the temperature and spot size derivatives obtained numerically. Note that due to the much higher sampling rate of the pyrometer, dT is an array, and due to the much lower sampling rate of the coaxial camera, dD is a scalar. We simply add the scalar to the array. If a higher sampling rate for the coaxial camera is used in the future, a nearest neighbor interpolation will be used to add the arrays.

Physics informed features are included into the network as additional inputs that have features extracted by the feature extraction network and are fused by the sensor fusion network as shown in Figure 3.3.

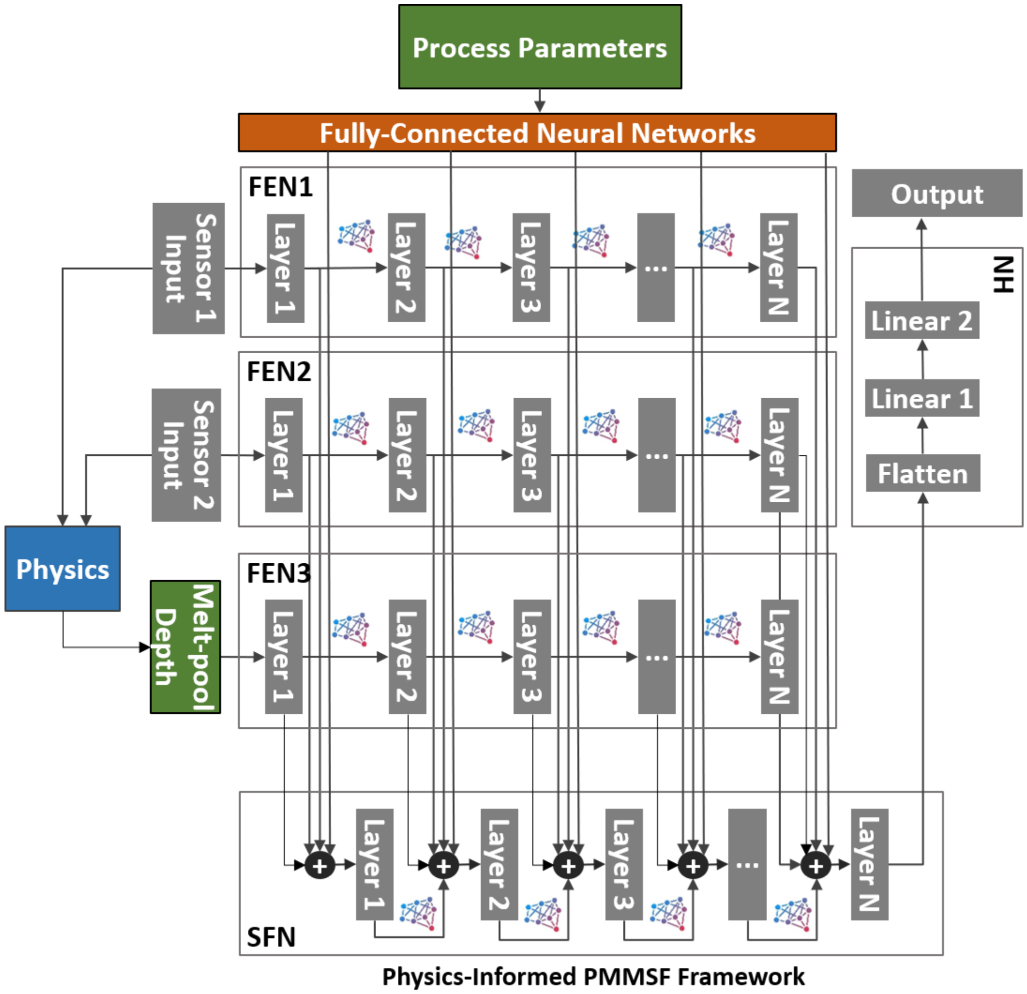

Figure 3.3. Network architecture for including physics-informed parameters. Process parameters inform each layer by weighing them with a fully connected neural network and SoftMax. Physics parameters are extracted from sensor signals and input through a separate FEN.

3.6 Process-Informed Features

Process-informed features represent the nominal operating regime of each sample. These features include laser power, scan speed, hatch spacing, layer height, nominal laser spot diameter, volumetric energy density, and areal energy density. The process parameters are not used as ground-truth defect labels; rather, they provide context for interpreting sensor signals. A thermal-emission value, melt-pool-image feature, or AE response may have a different meaning under different laser power or scan-speed conditions (King et al.; DebRoy et al.; Karniadakis et al.).

Process-informed inputs can be concatenated with learned sensor features or supplied to a weighting network that influences fusion. In the process-informed fusion configuration, the model is allowed to learn whether the reliability or importance of a sensor modality changes across process conditions. This is important because a sensor that is informative in a higher-energy-density condition may not contribute in the same way under a lower-energy-density condition, and vice versa (Baltrusaitis et al.; Satterlee et al. "Parallel Multi-Layer Sensor Fusion"; Karniadakis et al.).

3.7 Multi-Layer Context

The learning dataset includes current-layer and previous-layer context because pore formation may depend on both the immediate melt event and the local history of the material below it. Current-layer features describe the process signature associated with the queried location. Previous layer features can provide information about remelting, local surface condition, accumulated topography, thermal history, or powder-bed irregularities that may affect the current layer (King et al.; DebRoy et al.; Khairallah et al.).

The multi-layer formulation also supports the query-volume label definition used for registered XCT pores. A pore may occupy a finite height across several XCT slices and may intersect one or more process layers depending on its size and registration. The model therefore treats the current location as part of a local volume rather than as an isolated point sample. In this study, 5 layers are used. Additional layers account for defect movement such as pore formation in a layer below then remelting causing the pore to move to subsequent layers. When the stride is 1, pores only below the surface of the current layer are marked within that layer to better replicate layer by layer additive manufacturing. When a stride is greater than one is used, any intersection over the pore center is used. An example of pore labeling is shown in Figure 3.4.

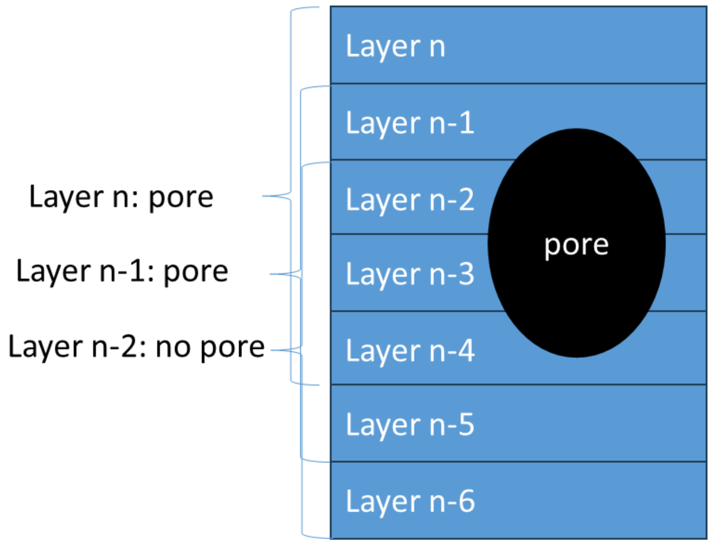

Figure 3.4. Layer-wise labeling and input scheme. A label is positive the pore center falls within the labeled layer volume

3.8 Fusion Comparisons and Sensor-Suite Optimization

The modeling approach evaluates individual sensors, reduced sensor combinations, full sensor fusion, input concatenation, head fusion, PMMSF, process-informed PMMSF, and physics-informed PMMSF. To keep fusion comparisons as close as possible, the number of layers and channels is kept the same between models as shown in Table 3.1. The fusion blocks are composed of 8 layers with progressively increasing channel sizes doubling each layer from 8 to 1024. The kernel size starts at 2 for the first layer and increases to 3 for each subsequent layer. The kernel stride starts at 1 for the first layer and increases to 2 for each subsequent layer. A dropout of 0.1 was used for each layer. The conv2D blocks have a kernel size of 3 and 5 layers with 16,32,32,32,64 channels. During training a batch size of 512 was used. Larger batch sizes were found to reduce early overfitting. This comparison is required because the best-performing sensor suite may not be the most practical configuration for deployment. A reduced sensor suite that preserves most of the predictive performance may be preferable if it lowers hardware complexity, integration difficulty, or data-processing burden (Everton et al.; Grasso and Colosimo; Satterlee et al. "Parallel Multi-Layer Sensor Fusion").

Table 3.1. Shared fusion model architecture. The block description here corresponds with those in the model architecture

Type	# Layers	Layer Channel	Dropout	kernel	stride
Fusion Block	8	8 to 1024	0.1	2,3,3,3,3,3,3,3	1,2,2,2,2,2,2,2
Conv2D	5	16,32,32,32,64	0	All 3	All 2

The input fusion architecture is presented in Figure 3.5. Each sensor is flattened and concatenated before being processed by the fusion blocks for feature extraction. The resulting features are compressed with a fully connected neural network with three blocks to a final size of 64 before being classified into two categories.

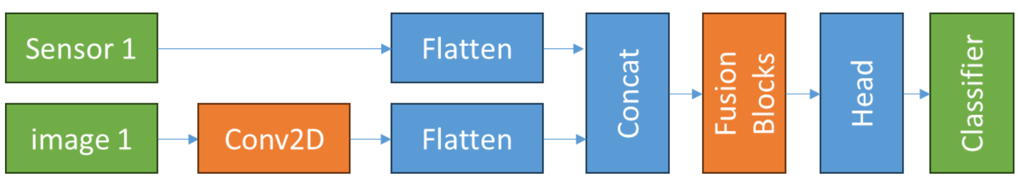

Figure 3.5. Architecture for input fusion methods. The fusion block is composed of 8 blocks and the Conv2D is composed of 5 2D convolutional blocks.

During head fusion, each sensor is processed independently by fusion blocks, and the outputs are concatenated before feeding into the head for feature fusion and finally classification as shown in Figure 3.6. Images first go through a conv2D prior to the fusion block for consistency.

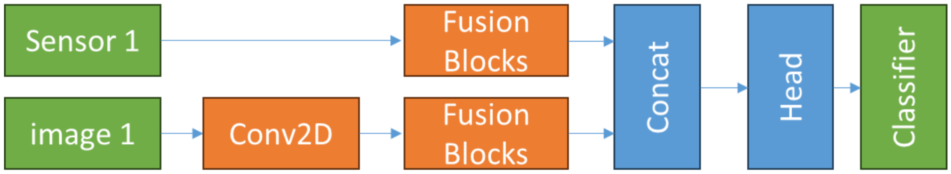

Figure 3.6. Architecture for head fusion method. The fusion block is composed of 8 blocks and the Conv2D is composed of 5 2D convolutional blocks.

Sensor-suite optimization is evaluated using the same localized label definitions so that differences in performance can be attributed to the sensor set or fusion strategy rather than to inconsistent labeling. The comparison also supports interpretation of whether pore-relevant information is concentrated in a single modality or distributed across thermal, acoustic, optical, topographic, and process-parameter inputs.

3.9 Training and Evaluation Strategy

For all reported comparisons, preprocessing, input normalization, training schedule, validation-threshold selection, and train/validation/test partitions were held fixed unless otherwise stated; only the sensor set, fusion architecture, or physics/process-feature configuration was changed.

Model performance is evaluated using classification metrics appropriate for sparse localized pore labels. The primary metrics are accuracy, F1 score, and AUC. Because the positive pore class is expected to be sparse relative to non-pore query volumes, class imbalance must be considered during sampling, training, and metric interpretation (Fawcett; Davis and Goadrich).

Accuracy is defined as the number of true predictions over the number of total predictions:

$$\text{Accuracy} = \frac{\text{TP+TN}}{\text{TP+TN+FP+FN}} \tag{3.17}$$

Where TP is the number of true positive cases, TN is the number of true negative cases, FP is the number of false positive cases, and FN is the number of false negative cases.

The F1 score is defined by the precision and the recall. It attempts to balance the two to provide a better metric for unbalanced datasets.

$$\text{Precision} = \frac{\text{TP}}{\text{TP+FP}} \tag{3.18}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP+FN}} \tag{3.19}$$

$$\text{F1} = 2\,\frac{\mathit{Precision}\cdot \mathit{Recall}}{\mathit{Precision}+\text{Recall}} = \frac{2\text{TP}}{2\text{TP+FN+FP}} \tag{3.20}$$

A binary class prediction is obtained by applying a decision threshold to the model output probability. Lowering the threshold increases the number of positive predictions and usually increases recall, but it can also increase false positives. Raising the threshold decreases the number of positive predictions and usually improves precision, but it can increase false negatives. ROC-AUC summarizes how well the model ranks positive samples above negative samples across thresholds; threshold-dependent metrics such as accuracy, precision, recall, and F1 depend on the selected threshold. An example ROC curve is shown in Figure 3.7.

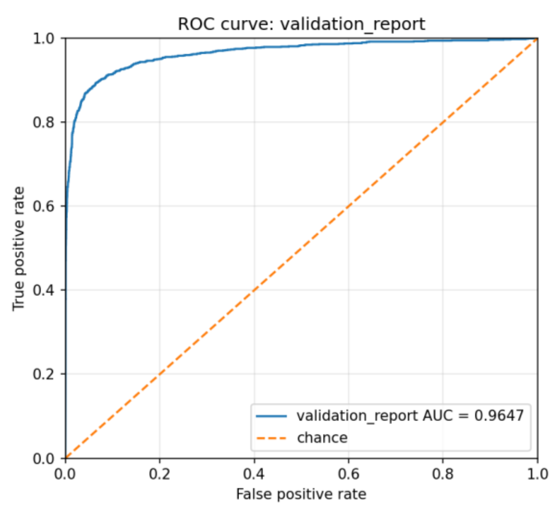

Figure 3.7. Example ROC curve for a pore prediction case. An intersection with the blue line indicates the results of the FPR and TPR.

Given a threshold value τ, the true positive rate and false positive rate are defined as

$$\mathrm{TPR}(\tau) = \frac{\mathrm{TP}(\tau)}{\mathrm{TP}(\tau)+\mathrm{FN}(\tau)} \tag{3.21}$$

$$\mathrm{FPR}(\tau) = \frac{\mathrm{FP}(\tau)}{\mathrm{FP}(\tau)+\mathrm{TN}(\tau)} \tag{3.22}$$

Then the AUC is ideally defined as the integral over these:

$$\mathrm{AUC} = \int_0^1 \mathrm{TPR}(\mathrm{FPR})\mathrm{dFPR} \tag{3.23}$$

$$\mathrm{AUC} \approx \sum_0^{n-1} \frac{\mathrm{TPR_{i+1}}+\mathrm{TPR_i}}{2}(\mathrm{FPR_{i+1}} - \mathrm{FPR_i}) \tag{3.24}$$

The closer AUC is to 1, the better the model generally performs.

The staged PMMSF training schedule is also used to make model comparisons more controlled. The same branch-pretraining, alternating feature/fusion training, branch-only fine-tuning, fusion-only fine-tuning, and joint fine-tuning sequence are used when comparing PMMSF configurations, process-informed PMMSF, and physics-informed PMMSF. The training schedule used is shown in Table 3.2. The best checkpoint is selected using the validation metric specified for the run, while the final test set is reserved for reporting.

Branch pretraining is performed to initially extract features from the sensor data independently. Only five epochs are used during pretraining to prevent strong features from overfitting the whole network. Alternating regularization uses two optimizers: one for the fusion network and one for the feature extraction networks. The training alternates between the two by updating one and freezing the weights of the other. In theory, this allows the features to align and be extracted by the fusion network. Fusion only fine-tuning trains only the fusion network while freezing the feature network weights, and joint fine-tuning utilizes one optimizer to train the whole network. The training regime is shown to be broken down in the loss in Figure 3.8. In previous studies it was found that jumping to joint training before pre-training and alternating regularization resulted in poor model performance. Therefore, the training procedure is as important as the network architecture.

Table 3.2. PMMSF Training Schedule

Fusion Schedule	Epochs	Learning Rate	Train FEN	Train SFN	Train Head	Alternate Freeze
pretraining	5	0.001	TRUE	FALSE	FALSE	FALSE
alternating_regularization	40	0.001	TRUE	TRUE	TRUE	TRUE
branches_only_refine	20	0.0009	TRUE	FALSE	TRUE	FALSE
fusion_only_refine	20	0.0008	FALSE	TRUE	TRUE	FALSE
joint_finetune	20	0.0006	TRUE	TRUE	TRUE	FALSE

All inputs are normalized based on the training set. During training data is augmented by applying noise and cropping. Up to 15% std noise can be applied and up to 15% crops can be applied. The probability of augmentation is decreased with a decay schedule of 5 epochs (Satterlee et al, "Automatic Detection"; Satterlee et al "Robust Image-Based") as shown in the Equation below:

$$p_{augment} = p_o \exp\left(-\frac{ln(2)}{\lambda} epoch\right) \qquad (3.25)$$

Where p_o is the initial probability, λ is the augmentation half-life, and epoch is the epoch number.

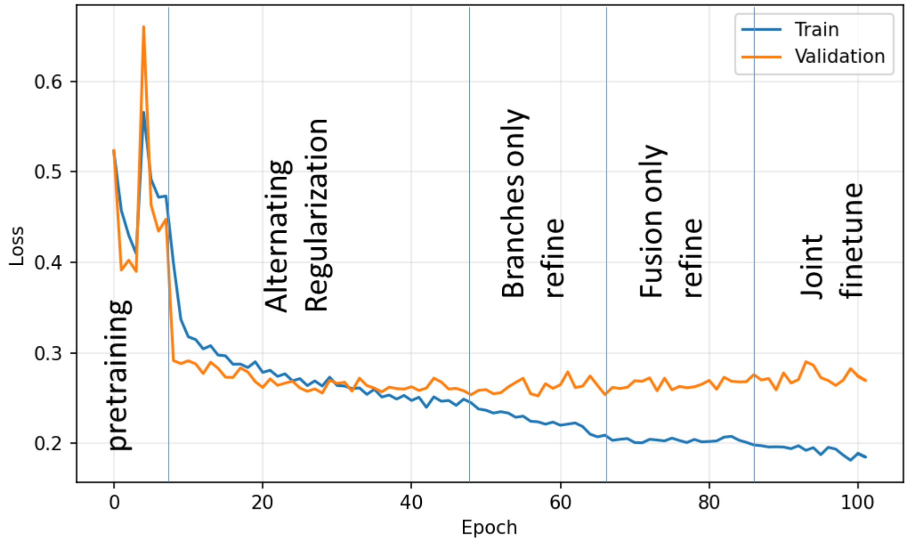

Figure 3.8. Loss graph with training schedule breakdown

The architecture with all sensor signals, process parameters, and extracted physics features is shown in Figure 3.9. This figure is generated by the custom code developed for this application called Fusion Studio. The input size of each sensor is driven by the slowest scan speed and sensor with highest sampling rate. The image features are first extracted via a 5-layer conv2d neural network before being linearized, then max pooling is applied to the linear features to match the 1D input requirements. The gate inputs are concatenated physics and process parameters. They are input into a 5 layer fully connected NN and then to a SoftMax to weigh each of the sensor signals. Weighing is applied via an addition operator and increases features that are more important depending on physics and process.

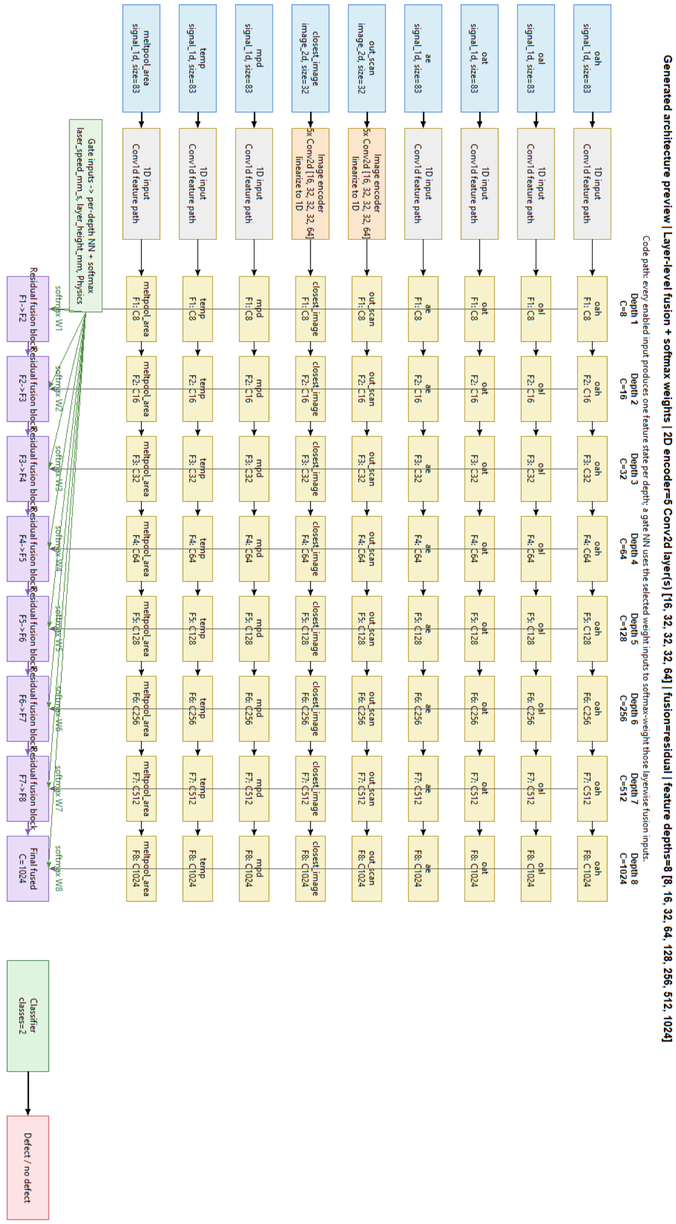

Figure 3.9. Generated architecture of employed PMMSF with all sensor signals, process parameters, and physics feature inputs.

The training, validation, and testing split uses a hybrid structure. The training is performed on the bottom 70% of 11 parts. Validation is performed on the top 30% of those same 11 parts. Testing is performed 100% of the held-out part. Of the 216 layers, only the top 166 layers were used for training. The training/validation/testing structure is visually represented in

Figure 3.10. The bottom of the part was designed for alignment and contained structural material. It contained additional defects that would otherwise be removed from post-processing. Since each part has a different combination of power and scan speed, the testing part is expected to better represent an independent trial. The query volume was fixed at 160 µm transverse × 300 µm scan direction × 150 µm build direction for all reported comparisons. This size was selected to balance localization precision, tolerance to residual registration error, and coverage of the largest accepted pores. Because each training case required up to approximately 10 hours, a full query-volume sensitivity study was not performed.

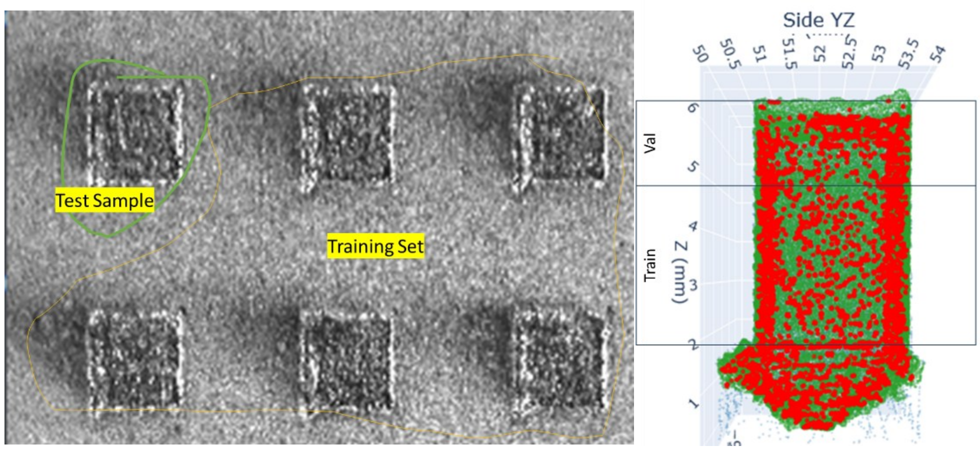

Figure 3.10. Visual representation of training and testing population. One part is held out for testing only (right), and the remaining parts are used for training and validation (left).

3.10 Chapter Summary

This chapter defines the feature-engineering and sensor-fusion approach used for localized LPBF pore prediction. The approach treats XCT-registered query-volume labels as the supervised target, constructs modality-specific sensors and feature inputs, includes multi-layer context, and evaluates PMMSF against simpler fusion baselines. Physics-informed and process-informed features are used to guide interpretation and fusion rather than to replace sensor data.

The next chapter describes the experimental system used to acquire the multimodal LPBF dataset.

Chapter 3 in part contains material as it appears in Parallel multi-layer sensor fusion for pipe leak detection using multi-sensors and machine learning. N. Satterlee, X. Zuo, C.-W. Lee, C.-W. Park, and J. S. Kang, Engineering Applications of Artificial Intelligence, vol. 153, p. 110923, 2025. The dissertation author was the primary investigator and author of this paper.

Chapter 3 in part contains material as it appears in Feature-informed machine learning for detecting material deformation and failure in aluminum pipes under bending load using acoustic emission sensors. X. Zuo, N. Satterlee, C.-W. Lee, I.-G. Choi, C.-W. Park, and J. S. Kang, *Materials & Design*, vol. 254, p. 114087, 2025. The dissertation author developed feature-informed network architecture and performed analysis on the results.

Chapter 4 Experimental System and Multimodal Sensor Acquisition

4.1 Overview

This chapter describes the experimental system used to generate the multimodal laser powder bed fusion dataset. The purpose of the experimental system was to collect in-situ sensor data during printing while preserving the process, spatial, temporal, and layer information needed for later synchronization to XCT-detected pore locations. The workflow included LPBF fabrication, multimodal sensor acquisition, layer tracking, after-spread and after-scan imaging, photometric-stereo heightmap generation, and structured data storage using the SMMVpy2 acquisition software.

The experimental system was designed to support localized porosity prediction rather than part-level defect classification. For this reason, the acquisition workflow required known process parameters, programmed geometry, scan-vector information, time-resolved sensor data, layer tracking, and organized data products suitable for XCT registration, localized label generation, feature extraction, and PMMSF model training.

This chapter is limited to the experimental platform, process matrix, sensor suite, acquisition software, and raw or derived data products generated during printing. Coordinate synchronization, XCT registration, pore segmentation, localized label generation, and machine-learning dataset construction are described in later chapters.

4.2 LPBF Platform, Material, Build Geometry, and Process Matrix

The multimodal dataset was generated using an Aconity Mini LPBF system equipped with a 1070 nm laser shown in Figure 4.1. The material was m4p 316L.03 stainless steel powder. The coupons were fabricated on a 316L stainless steel build plate under argon shielding gas, with the oxygen concentration maintained below 300 ppm O_2 during printing and the fume extraction pump at 50% duty cycle. This material and process window were selected to support localized analysis of XCT-detected porosity in nominally dense LPBF 316L stainless steel.

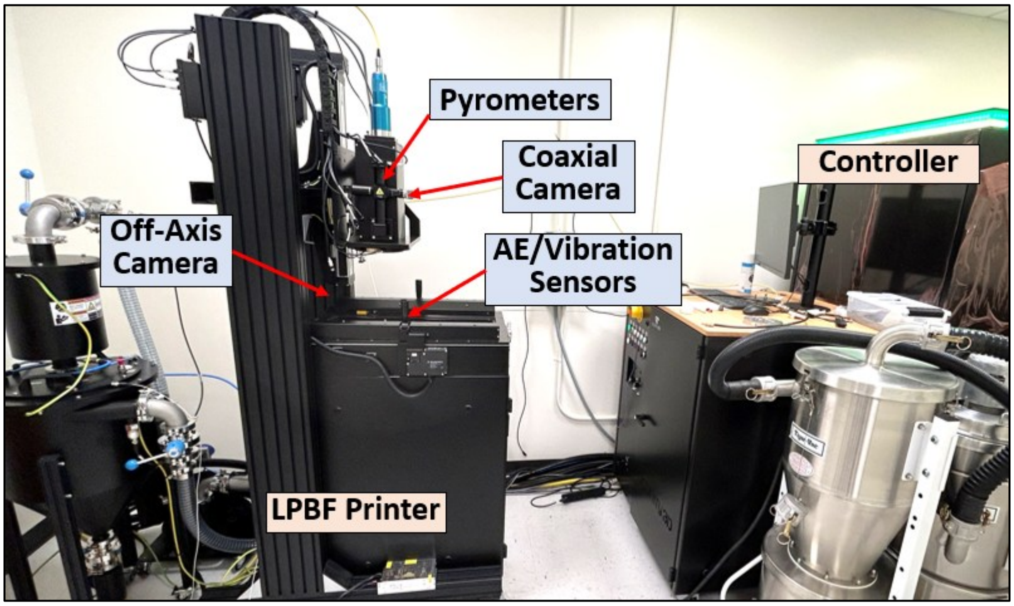

Figure 4.1. Aconity Mini LPBF machine and installed sensor system

Twelve coupons were produced from six programmed CLI geometries with print location shown in Figure 4.2. Each CLI file was reused across two process-parameter groups, denoted print 1 and print 2. Print 1 used a higher volumetric energy-density range, while print 2 used a lower volumetric energy-density range as shown in Table 4.1. Both groups were treated as nominally dense process conditions, and neither group was assumed to represent a fully defective or non-defective class. The supervised labels used in later chapters were generated from XCT-detected porosity rather than process group assignment.

The programmed coupon envelope was approximately 4 mm × 4 mm × 6.48 mm. This programmed envelope was used for geometric synchronization because the exterior geometry was later used during XCT-to-process registration. Each CLI file contained 216 layers with both contour and hatch scan vectors. Hatch rotation occurred across layers, with scan directions appearing approximately at 0°, 45°, 90°, and 135°. The layer height was 30 µm, hatch spacing was 80 µm, and nominal laser spot diameter was 80 µm.

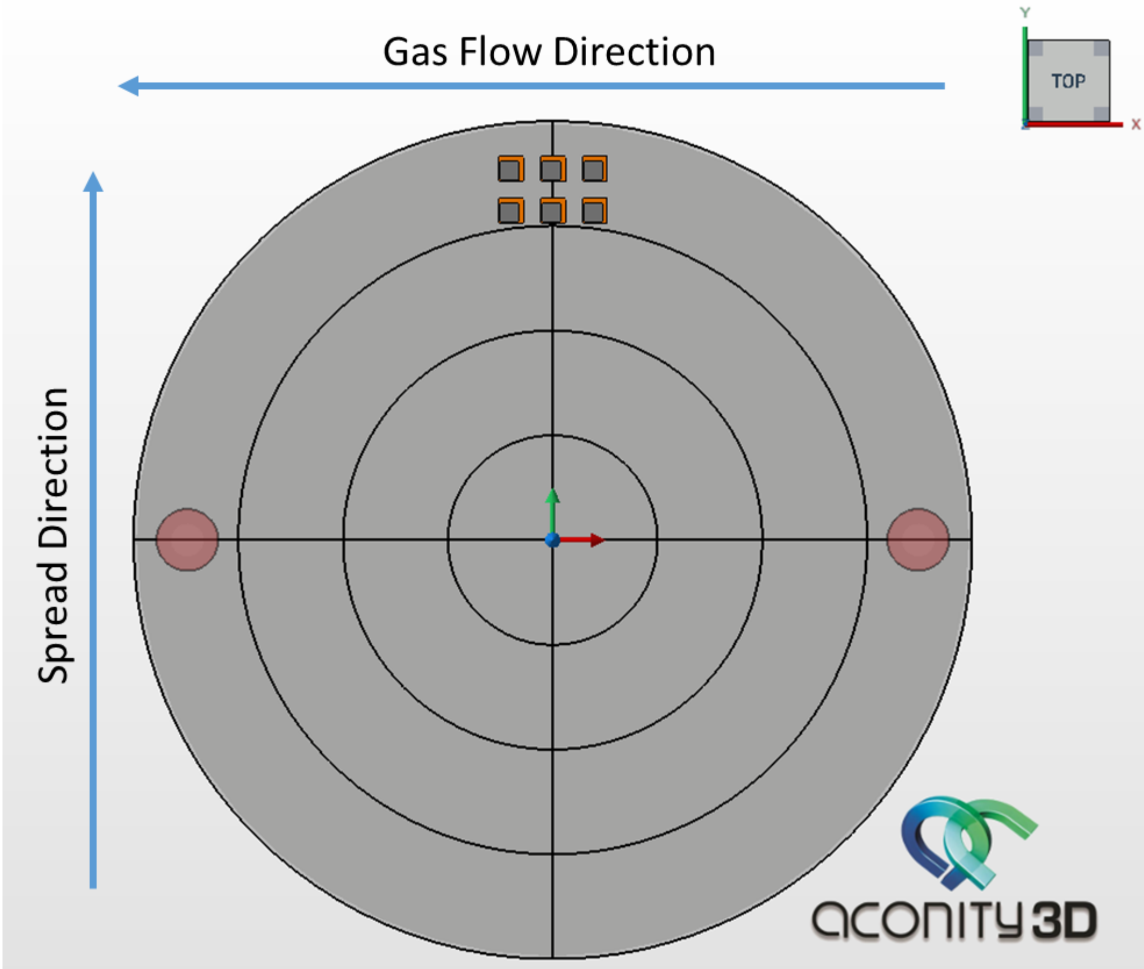

Figure 4.2. Build layout for six LPBF coupons showing coupon placement on the build plate, gas-flow direction, and spread direction

The process matrix is summarized in Table 4.1. Volumetric energy density, P_E, and areal energy density, P_C, were calculated as

$$P_E = \frac{P}{vht} \tag{4.1}$$

$$P_C = \frac{P}{vd} \tag{4.2}$$

where (P) is laser power, (v) is scan speed, (h) is hatch spacing, (t) is layer height, and (d) is nominal laser spot diameter.

Table 4.1. LPBF process-parameter matrix for the 316L stainless steel coupons.

Print	Part	Laser Power (W)	Scan Speed (mm/s)	Hatch (µm)	Spot (µm)	(P_E) (J/mm³)	(P_C) (J/mm²)	XCT Voxel Size (mm)
1	1	190	1100	80	80	71.97	2.159	0.0076
1	2	170	1000	80	80	70.83	2.125	0.0076
1	3	155	900	80	80	71.76	2.153	0.0076
1	4	135	800	80	80	70.31	2.109	0.0076
1	5	120	700	80	80	71.43	2.143	0.0076
1	6	100	600	80	80	69.44	2.083	0.0076
2	7	130	1100	80	80	49.24	1.477	0.0076
2	8	120	1000	80	80	50	1.5	0.0076
2	9	110	900	80	80	50.93	1.528	0.0076
2	10	95	800	80	80	49.48	1.484	0.0076
2	11	85	700	80	80	50.6	1.518	0.0076
2	12	75	600	80	80	52.08	1.563	0.0076

The programmed geometry and scan-vector content are summarized in Table 4.2. The CLI files contained both contour and hatch content as shown in Figure 4.3; therefore, the scan strategy was not treated as hatch-only. This distinction is important because contour vectors, part edges, and scan-vector transitions can produce sensor signatures that differ from interior hatch regions (King et al.; DebRoy et al.).

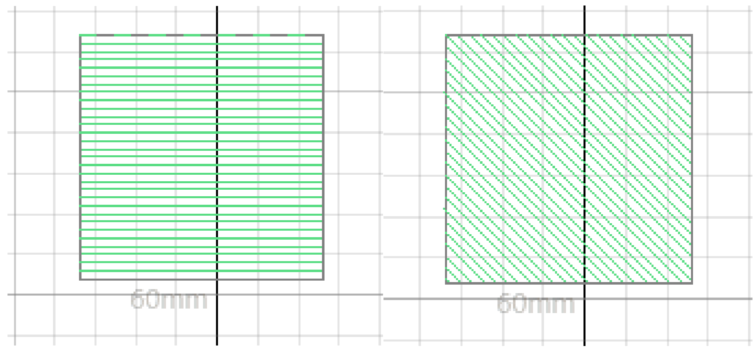

Figure 4.3. CLI-defined coupon geometry and scan strategy. The programmed coupon envelope, contour vectors, hatch vectors, and representative hatch are shown for 2 layers (layer 75 left and layer 76 right).

Table 4.2. CLI file and programmed-geometry summary.

CLI File	X Envelope (mm)	Y Envelope (mm)	Z Envelope (mm)	Layers	Polyline Records	Hatch Segments
P1, P7	-9 to -5	60 to 64	0.030 to 6.5	216	1,794	9,044
P2, P8	-2 to 2	60 to 64	0.030 to 6.5	216	1,794	8,996
P3, P9	5 to 9	60 to 64	0.030 to 6.5	216	1,794	9,044
P4, P10	-9 to -5	53 to 57	0.030 to 6.5	216	1,794	9,086
P5, P11	-2 to 2	53 to 57	0.030 to 6.5	216	1,794	9,038
P6, P12	5 to 9	53 to 57	0.030 to 6.5	216	1,794	9,086

4.3 XCT system

Post-build X-ray computed tomography was performed to obtain three-dimensional pore ground truth for the printed LPBF coupons. The samples were scanned using a North Star Imaging X5000 system equipped with an X-Ray WorX 225 kV microfocus source and a PerkinElmer XRD 1621 flat-panel detector with a 200 µm pixel pitch as shown in Table 4.3. The XCT scans were used to reconstruct coupon volumes for subsequent part isolation, pore segmentation, XCT-to-process registration, and localized pore-label generation.

The scans were acquired using tube voltages of 220 kV and tube currents between 28 and 35 µA. The microfocus focal spot size ranged from approximately 6.2 to 7.7 µm. Each scan used 1,440 to 2,160 projections with 6 to 8 frame averages per projection. The source-to-object distance was approximately 30 to 35 mm, and the source-to-detector distance ranged from 856.5 to 996.1 mm. These acquisition conditions produced an effective voxel size of approximately 6.3 to 8.0 µm. The acquisition time was approximately 3 hours per scan. The list of XCT parameters is shown in Table 4.3

Table 4.3. XCT acquisition parameters

Parameter	Value
XCT system	North Star Imaging X5000
X-ray source	X-Ray WorX 225 kV microfocus source
Detector	PerkinElmer XRD 1621 flat-panel detector
Detector pixel pitch	200 µm
Tube voltage	220 kV
Tube current	28–35 µA
Focal spot size	6.2–7.7 µm
Frame averaging	6–8 frames
Number of projections	1,440–2,160
Acquisition time	Approximately 3 h per scan
Source-to-object distance	30–35 mm
Source-to-detector distance	856.5–996.1 mm
Effective voxel size	Approximately 6.3–8.0 µm

4.4 Multimodal Sensor Suite

The experimental system integrated machine process data, pyrometer signals, acoustic emission, a laser-trigger signal, coaxial melt-pool imaging, off-axis layer imaging, and photometric-stereo heightmap generation. The sensor suite was selected to provide complementary observability of the LPBF process. Machine-position and laser-process signals provided coordinate references used to map sensor data to the programmed build geometry. The separate laser-trigger signal was recorded with AE on the NI myDAQ to support acoustic synchronization. Pyrometer channels provided relative thermal-emission signatures as shown in Figure 4.4. Acoustic emission provided high-frequency process response transmitted through the build hardware as shown in Figure 4.7. Coaxial imaging provided melt-pool geometry and intensity information as shown in Figure 4.5. Off-axis imaging and photometric stereo provided layer-wise surface and powder-bed context as shown in Figure 4.6 (Everton et al.; Grasso and Colosimo; Tapia and Elwany). The layout of the sensors and data collection devices is shown in Figure 4.8. An overview of the hardware, acquisition settings, and optical or mounting configuration is shown in Table 4.4.

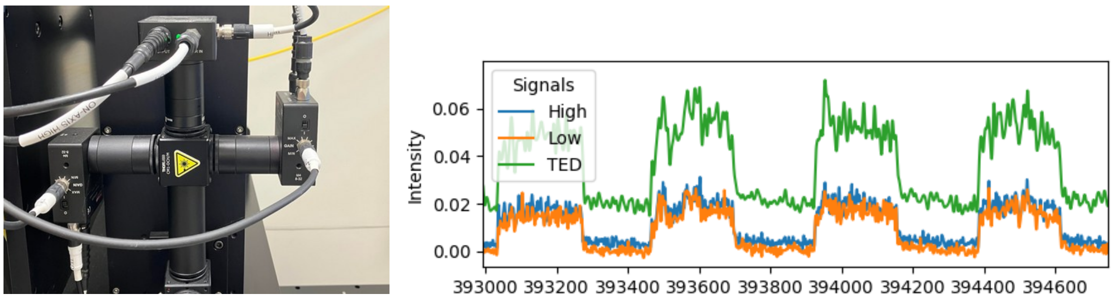

Figure 4.4. Pyrometers and sample signal output at 50ks/s and 25ks/s

Figure 4.5. Coaxial camera and sample signal output with 8um pixel size

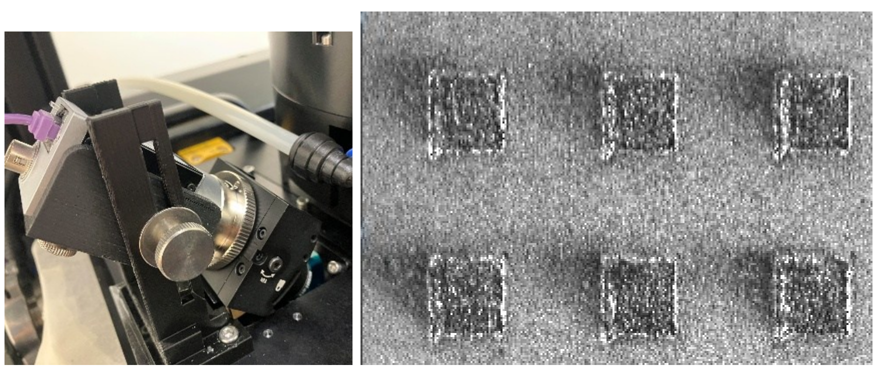

Figure 4.6. Off axis camera and sample output of post-scanned powder bed with 5um pixel size

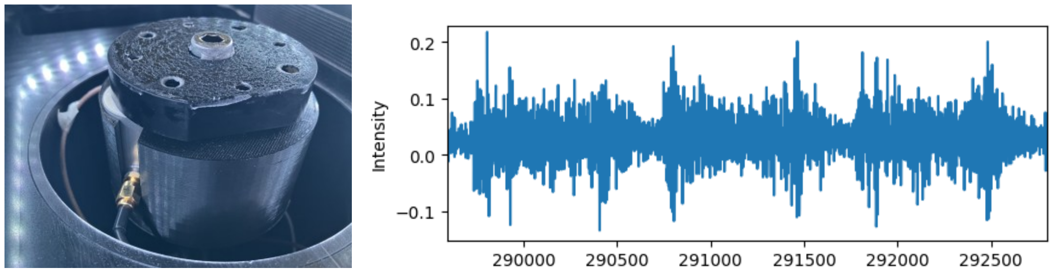

Figure 4.7. AE sensor with sample signal output at 200ks/s

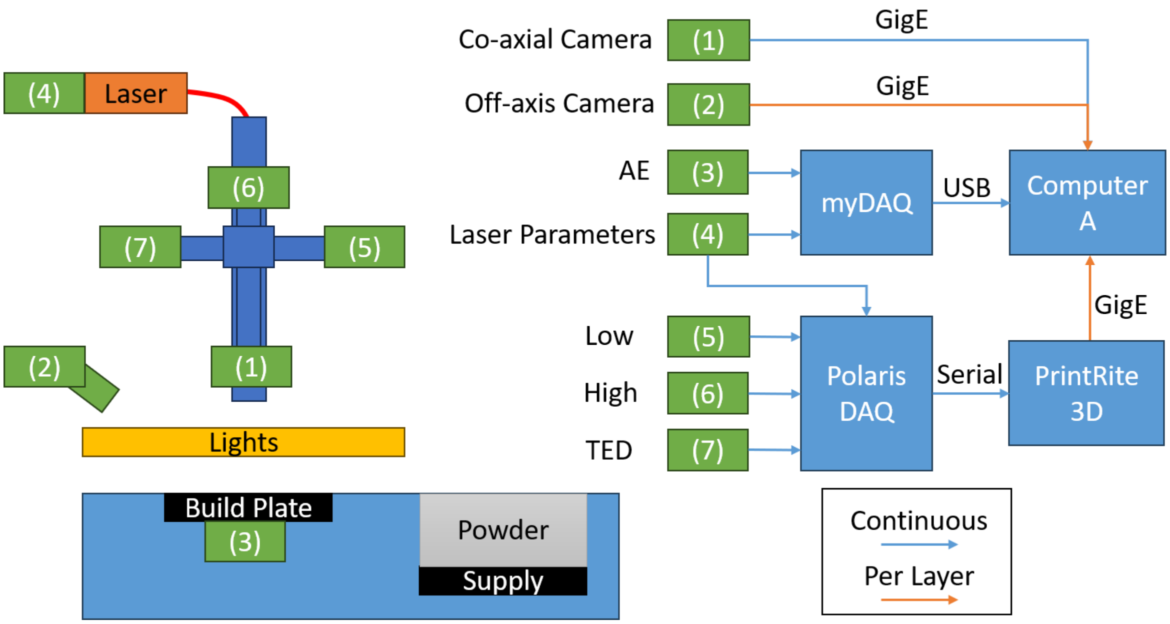

Figure 4.8. Experimental sensor layout for the multimodal LPBF monitoring system including the Aconity Mini build chamber, coaxial optical path, PrintRite3D/Polaris process-data acquisition, off-axis camera and lighting, NI myDAQ acquisition for acoustic emission and laser trigger, and the AE sensor mounted below the build plate through the build-plate/mounting-bracket mechanical chain.

Table 4.4. Summary of the hardware, acquisition settings, optical or mounting configuration, and primary data product for each sensing modality.

Modality	Hardware	Acquisition Setting	Optics / Mounting	Primary Data Product
Laser position/process	Sigma PrintRite3D with Polaris Motion / Polaris DAQ architecture	Machine/process acquisition; position and process signals	Machine-integrated monitoring system	XY data, laser modulation/drive, laser power, raw HDF5/process files
Pyrometer High	Thorlabs APD440A-SP1	50 kS/s; nominal 700 nm channel	PrintRite3D optical path with filters/beamsplitter	High-channel thermal-emission time series
Pyrometer Low	Thorlabs APD440A-SP1	50 kS/s; nominal 680 nm channel	PrintRite3D optical path with filters/beamsplitter	Low-channel thermal-emission time series
TED	Thorlabs APD440A-SP1	25 kS/s	PrintRite3D TED optical path	Thermal energy density channel
AE + laser trigger	AES-R60 + laser-trigger channel through NI myDAQ	200 kS/s; two analog input channels	AE grease-coupled to mounting bracket below build plate; laser trigger recorded on same DAQ	Synchronized AE and laser-trigger time series
Coaxial imaging	Teledyne DALSA Genie Nano M640-NIR	8000 fps; cropped to 64 × 24 px; exposure 40μs; 8μm/pixel	Kowa LM100FC24M lens; VA-LADAP-LCM-4xF-150 focal-length extender	Coaxial melt-pool image sequence
Off-axis imaging and photometric stereo	Basler acA5472-5gm	Layer-wise after-spread and after-scan capture; exposure 80000μs; ~5μm/pixel; eight directional illumination images	Opto Engineering MCSM1-01X lens; VA-LADAP-LCM-2xF-150 focal-length extender; directional illumination geometry	Directional off-axis image stacks, all-light display image, and derived relative heightmaps

Together, these systems produced synchronized process, acoustic, thermal-emission, melt-pool image, and layer-surface data for later comparison with XCT-detected porosity. The data were preserved by layer, time, and process coordinate so that localized process signatures could be evaluated without assuming a single pore-formation mechanism.

4.5 SMMVpy2 Data Acquisition Software

SMMVpy2 was used to coordinate multimodal data acquisition during the build. The software provided a PyQt-based interface for data collection, calibration, printer control, off-axis image capture, heightmap generation, post-processing, and data synchronization. During printing, the software initialized acquisition from the NI analog input device, the Teledyne DALSA coaxial camera, and the Basler off-axis camera while coordinating layer tracking and printer wait/not-wait commands. The commands were integrated into the system via digital inputs from an ESP32 microcontroller. A custom board was designed to regulate the input and output voltages and facilitate communication between the machine and the microcontroller. The microcontroller was then controlled by SMMVpy2 software which periodically reads values from the microcontroller and sends updated values to it. A high-level overview of the SMMVpy2 software is shown in Figure 4.9.

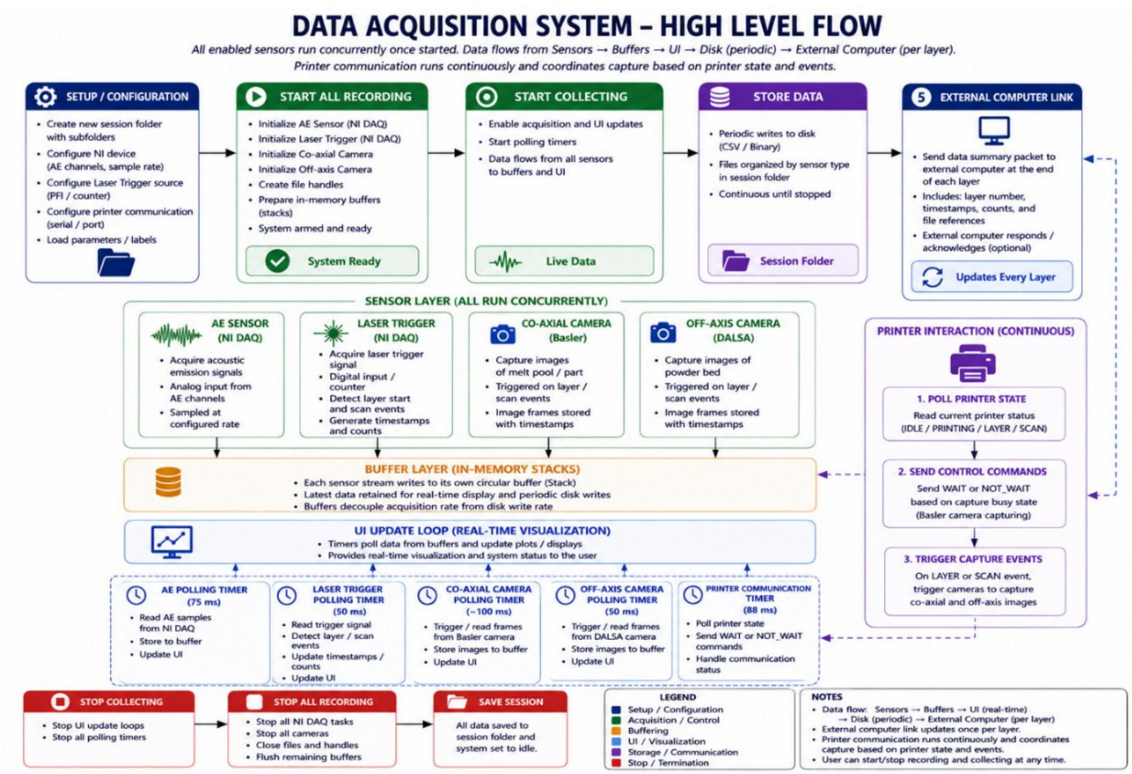

Figure 4.9. SMMVpy2 acquisition workflow including the printer-state communication, layer-edge detection, scan-edge detection, NI acquisition, coaxial image acquisition, Basler off-axis image acquisition, photometric-stereo processing, and structured data storage.

The acquisition workflow created a structured directory for each print and session. Data were organized into modality-specific subfolders, including NI device data, coaxial camera data, Basler off-axis images, and pyrometer/process data. This structure preserved the raw data while supporting later conversion into synchronized process-coordinate datasets.

Printer-state communication was used to detect layer and scan transitions. The software returned one state for a layer edge and another state for a scan edge. These state values were used to coordinate acquisition and to pause or resume the printer when off-axis imaging was required. On a new layer, the software also downloaded the latest pyrometer file from the remote host into the local pyrometer-data folder.

The SMMVpy2 workflow served as the coordination layer between the LPBF printer, machine process-data system, high-speed coaxial imaging, layer-wise off-axis imaging, AE acquisition, laser-trigger acquisition, and post-processing pipeline.

4.6 PrintRite3D Process and Pyrometer Data

PrintRite3D/Sigma Labs was used to acquire laser position, process, and pyrometer data through Polaris Motion / Polaris DAQ architecture. The acquired signals included laser position, laser modulation or drive, laser power, and pyrometer/process channels. These data provided the primary process-coordinate reference for later mapping of in-situ signals to the machine coordinate system. Conversion from laser coordinates (X-Position, Y-Position) to machine coordinates (X-Machine, Y-Machine) was performed by mapping laser locations to known positions specified in the CLI file.

Three Thorlabs APD440A-SP1 channels were acquired through the PrintRite3D optical path. The channels were designated High, Low, and TED. The High channel corresponded approximately to a 700 nm pyrometer channel, the Low channel corresponded approximately to a 680 nm pyrometer channel, and TED represented the thermal energy density channel. The High and Low were acquired at 50 KS/s and the TED was acquired at 25 KS/s. These signals are treated as relative thermal-emission and process-response signals rather than calibrated absolute temperature measurements because absolute thermal interpretation would require calibration, emissivity assumptions, and optical-path characterization beyond the scope of this acquisition workflow.

The PrintRite3D process data supported coordinate synchronization and process-informed feature extraction. The final supervised labels were not taken from PrintRite3D anomaly flags or metric images; labels were generated from XCT-detected pore morphology after registration to the process coordinate system.

4.7 Acoustic Emission and Laser-Trigger Acquisition

Acoustic emission was measured using an IDK-AES-R60 sensor sampled at 200 kS/s using the NI myDAQ. The frequency range is from 40-350kHz with resonant frequency of 60 +/- 15kHz. It was selected because it contained the widest band within the Nyquist of the sampling range. The AE sensor was grease-coupled to the mounting bracket below the build plate, giving a mechanical transmission path from the build plate to the mounting bracket and then to the AE sensor. This configuration allowed process-induced acoustic and structural response to be measured outside the build chamber while maintaining mechanical coupling to the printed part and build platform.

The laser-trigger signal was acquired on the same NI myDAQ as the AE signal. Both analog input channels of the NI myDAQ were used: one channel for AE and one channel for the laser trigger. This configuration preserved a direct timing relationship between acoustic events and laser operation, which was necessary because localized process instabilities may occur over short time scales. Recording the trigger and AE together allowed AE features to be aligned to laser activity in later synchronization steps (Gorgannejad et al.; Hocker et al.).

The AE signal was treated as a high-frequency process-response signal rather than a direct image of pore formation. It may contain information related to melt-pool dynamics, spatter, keyhole-related instability, powder interaction, structural vibration, or other transient events (Gorgannejad et al.; Everton et al.).

4.8 Coaxial Melt-Pool Imaging

Coaxial melt-pool images were acquired using a Teledyne DALSA Genie Nano M640-NIR camera. The camera was operated at 8000 frames per second using a cropped 64 × 24-pixel region and an exposure setting of 40µs. The optical setup included a Kowa LM100FC24M lens and a VA-LADAP-LCM-4xF-150 focal-length extender resulting in a pixel resolution of about 8µm. The camera was integrated into the coaxial optical path to observe the process zone during laser scanning. A 400-700nm cutoff filter was applied to the camera to reject reflected laser light. To reduce memory requirements, the meltpool image was cropped to 32 × 24-pixel prior to machine learning training.

The coaxial camera provided image sequences of melt-pool intensity and projected melt-pool geometry. Melt-pool area, shape, intensity distribution, and temporal fluctuations can reflect local changes in energy coupling and melt-pool stability. However, the coaxial image stream measures surface-projected optical response and does not directly measure subsurface keyhole depth or pore formation. Coaxial features were therefore treated as indirect process signatures (Huang et al.; Zhao et al.; Gorgannejad et al.).

The SMMVpy2 coaxial acquisition module used high-speed Mono8 continuous acquisition with chunked writing, status and error logging, and restart handling. Coaxial image chunks were saved by layer and included image arrays, timestamps, frame-rate information, frame shape, data type, frame count, layer index, and start and end times.

4.9 Off-Axis Imaging and Photometric-Stereo Heightmaps

Layer-wise off-axis images were acquired using a Basler acA5472-5gm camera with an Opto Engineering MCSM1-01X lens and a VA-LADAP-LCM-2xF-150 focal-length extender resulting in a pixel resolution of 5µm. Images were captured after spreading and after scanning for each layer using an exposure setting of 80,000µs. Eight directional illumination images were captured for photometric-stereo reconstruction. An example of a four-directional-light case is shown in Figure 4.10. An additional all-light image was captured for real-time layer-wise display but was not used as a training input.

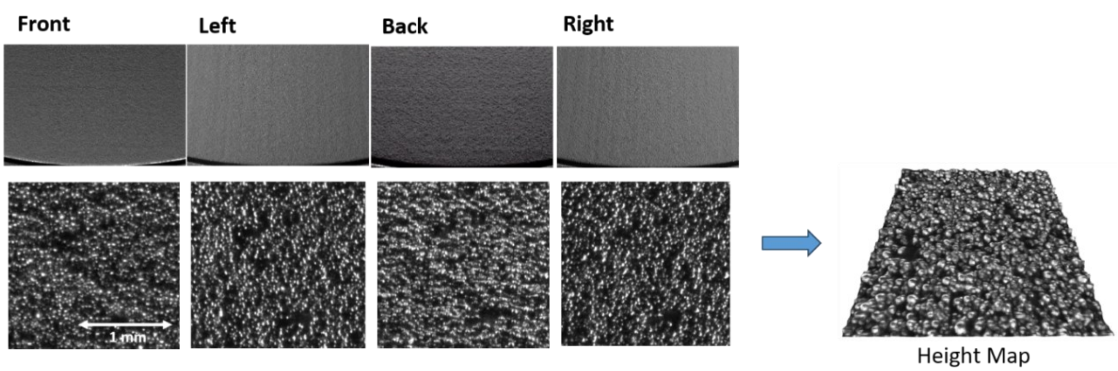

Figure 4.10. Off-axis image capture example with generated heightmap.

Off-axis imaging provided powder-bed and layer-surface context that could not be observed directly from coaxial melt-pool imaging or one-dimensional process signals. After-spread images described the powder-bed condition before laser exposure, while after-scan images described the solidified surface after processing. These measurements are relevant because powder-bed irregularities, surface topography, recoating effects, and previous-layer conditions may influence pore formation in subsequent layers (Everton et al.; Grasso and Colosimo; Woodham).

The photometric-stereo workflow included four stages shown in Figure 4.11. First, pixel calibration was performed before the print by capturing a long-exposure image while known laser locations were exposed. The laser points were localized in the image, mapped to laser position, and saved as a transform for later image capture. Second, image capture was performed before and after spreading or scanning by pausing the print, starting the LED sequence, capturing directional images, and saving the image set. Third, heightmaps were generated by normalizing images, defining lighting positions, estimating surface normals and albedo, creating a normal map, and integrating the normals. Fourth, post-processing and calibration were applied through high-pass filtering, coordinate transformation, localization to the laser position, and cropping. The mathematical reconstruction used for the heightmap generation is summarized below.

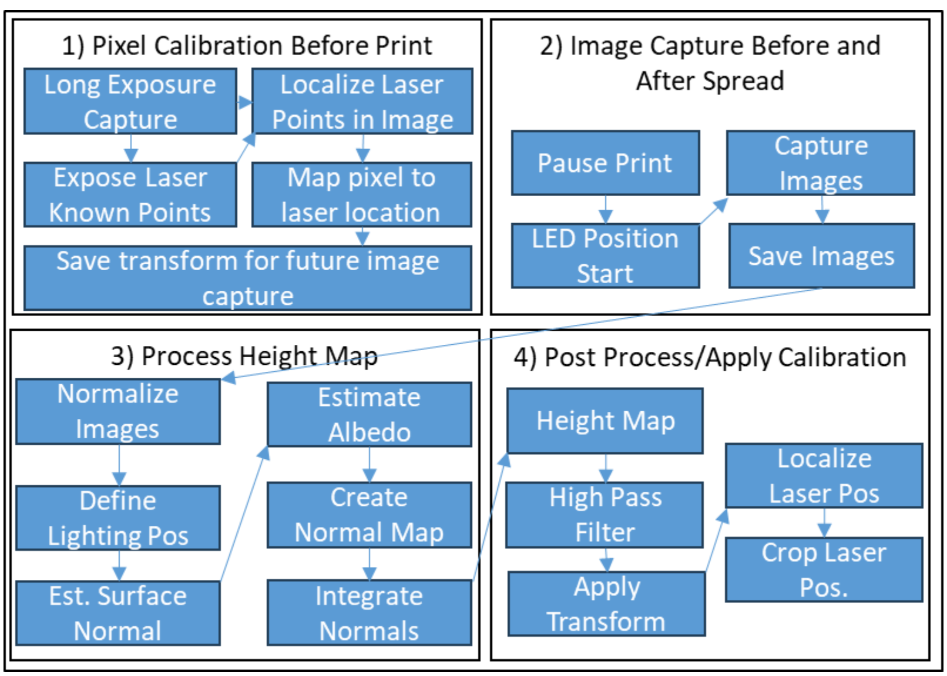

Figure 4.11. Photometric-stereo workflow used for off-axis layer-surface reconstruction, including pre-print pixel calibration, image capture before and after spreading/scanning, height-map generation, and post-processing through filtering, coordinate transformation, localization, and cropping.

Under the calibrated Lambertian photometric-stereo model introduced by Woodham, each pixel is imaged from a fixed camera position under multiple known illumination directions. For pixel location u and illumination direction i, the measured intensity is modeled as Equation 4.3 (Woodham).

$$I_i(\mathrm{u}) = \rho(\mathrm{u})\, \mathrm{s}_i^T \mathrm{n}(\mathrm{u}) + \epsilon_i(\mathrm{u}) \tag{4.3}$$

For q directional illuminations, the pixel intensities are stacked into the vector I(u), the lighting directions are assembled into the q by 3 matrix S, and the scaled-normal vector g(u) is defined as the product of albedo and the unit surface normal. This gives the linear system in Equation 4.4.

$$\mathrm{I}(\mathrm{u}) = \mathrm{S}\,\mathrm{g}(\mathrm{u}) + \epsilon(\mathrm{u}), \qquad \mathrm{g}(\mathrm{u}) = \rho(\mathrm{u})\,\mathrm{n}(\mathrm{u}) \tag{4.4}$$

When at least three non-coplanar illumination directions are available, the scaled normal is estimated by least squares using Equation 4.5. The albedo and unit surface normal are then recovered using Equation 4.6.

$$\hat{\mathrm{g}}(\mathrm{u}) = (\mathrm{S}^T\mathrm{S})^{-1}\mathrm{S}^T\mathrm{I}(\mathrm{u}) \tag{4.5}$$

$$\hat{\rho}(\mathrm{u}) = \|\hat{\mathrm{g}}(\mathrm{u})\|_2, \qquad \hat{\mathrm{n}}(\mathrm{u}) = \frac{\hat{\mathrm{g}}(\mathrm{u})}{\|\hat{\mathrm{g}}(\mathrm{u})\|_2} \tag{4.6}$$

The recovered normal field is converted to surface gradients using Equation 4.7 and integrated to obtain a relative heightmap. In the least-squares normal-integration formulation, the height field z satisfies the Poisson form in Equation 4.8; because measured normal fields can be noisy or non-integrable, the resulting heightmap is interpreted as a relative topographic feature unless separately validated (Frankot and Chellappa).

$$p(\mathrm{u}) = \frac{\partial z}{\partial x} = -\frac{n_x(\mathrm{u})}{n_z(\mathrm{u})}, \qquad q(\mathrm{u}) = \frac{\partial z}{\partial y} = -\frac{n_y(\mathrm{u})}{n_z(\mathrm{u})} \tag{4.7}$$

$$\nabla^2 z = \frac{\partial p}{\partial x} + \frac{\partial q}{\partial y} \tag{4.8}$$

In this workflow, the photometric-stereo equations provide the basis for estimating surface normals, albedo, and a relative height representation from the eight directional illumination images. Examples of surface normals, albedo, and a relative height representation are shown in Figure 4.12. The final registered heightmap is then filtered, transformed into the machine-coordinate frame, localized to the laser position, and cropped for downstream feature extraction. The heightmap image is 16-bit so high frequency powder bed topography is represented even if not easily observed due to the large low-frequency height differences.

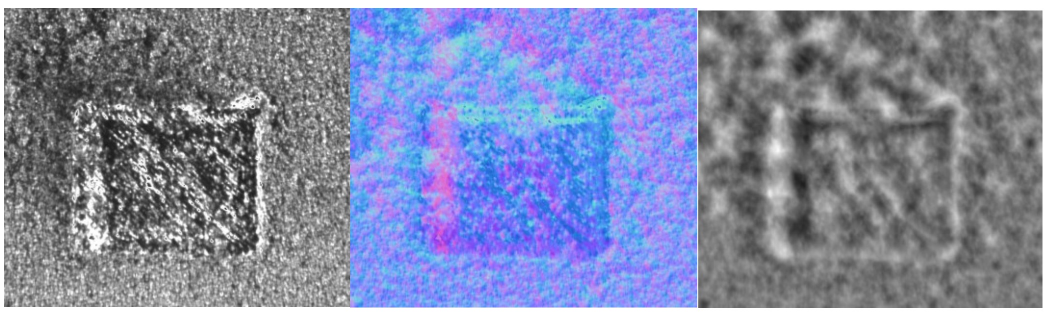

Figure 4.12. Estimated surface normals, albedo, and a relative height representation from the eight directional illumination images.

Photometric stereo provides surface-orientation and relative-height information from images acquired under different illumination directions. In this dissertation, the resulting heightmaps were used as relative topographic inputs rather than absolute metrology because the learning framework required localized surface-condition features registered to the build coordinate.

4.10 Layer Tracking and Printer Coordination

Layer tracking was required to match signals of the same layer in post processing. Printer-state communication in SMMVpy2 detected pulse signals from the printer that include layer edges and scan edges, allowing sensor data and image captures to be organized by layer. This layer information was later combined with CLI scan vectors, machine-position data, and XCT registration to generate localized learning samples.

The off-axis imaging workflow required coordination with printer motion. During selected capture events, the printer could be paused, the illumination sequence executed, images acquired, and the printer resumed. This coordination allowed after-spread and after-scan surface images to be captured without treating the off-axis camera as a continuous high-speed sensor. In contrast, AE, laser trigger, pyrometer channels, and coaxial imaging were acquired as time-resolved process signals during laser operation.

Layer tracking also supported previous-layer context in later dataset construction. Because pores may be influenced by the current scan, previous-layer surface condition, local part geometry, and neighboring process history, preserving layer identity was necessary for constructing multi-layer inputs to the PMMSF model (King et al.; DebRoy et al.; Khairallah et al.).

4.11 Data Organization

The acquisition workflow produced raw data, processed image products, calibration information, and intermediate files for later synchronization. These products are summarized in Table 4.5. The table separates data generated during printing from data generated during later post-processing because supervised-learning labels were produced only after XCT segmentation and registration.

Table 4.5. Raw, derived, and downstream data products used in the dissertation workflow.

Data Product	Source	Rate / Frequency	Format / Organization
CLI scan files	Aconity build preparation	One per programmed geometry	ASCII CLI; layer-wise contour and hatch vectors
PrintRite3D raw process data	PrintRite3D / Polaris Motion DAQ	Machine/process acquisition	Raw HDF5 time-series data
PrintRite3D metric image stacks	PrintRite3D analytics	Layer-wise metric data	Image-stack metric data
High pyrometer signal	APD440A-SP1 / PrintRite3D	50 kS/s	Time series
Low pyrometer signal	APD440A-SP1 / PrintRite3D	50 kS/s	Time series
TED signal	APD440A-SP1 / PrintRite3D	25 kS/s	Time series / metric channel
AE signal	AES-R60 / NI myDAQ	200 kS/s	Time series
Laser-trigger signal	Laser trigger / NI myDAQ	200 kS/s, acquired with AE	Time series
Coaxial melt-pool images	Teledyne DALSA Genie Nano M640-NIR	8000 fps	64 × 24 px cropped image stream
After-spread off-axis images	Basler acA5472-5gm	Once per layer	Eight directional images plus all-light display image
After-scan off-axis images	Basler acA5472-5gm	Once per layer	Eight directional images plus all-light display image
Photometric-stereo heightmaps	Directional off-axis image stack	Once per captured image set	Relative heightmap
Calibration images	Off-axis camera and laser marks	Calibration event	Image files and transform data
Machine-coordinate HDF5 data	Post-processing pipeline	One processed build dataset	HDF5 with process fields and machine coordinates
XCT volumes	XCT scan	One scan per 4 parts	3D reconstructed volume
Registered pore tables	XCT segmentation and registration	One table per registered part/build	HDF5 pore morphology table

The tensor size for each sensor is also dependent on the sensor type. The list of tensors for each sensor and typical values is shown in Table 4.6. Samples for each signal are captured contiguously for each layer and stacked into channels. This means that images are not just for the top layer. Images are stacked from previous layer and used as input for the current layer.

Table 4.6. Typical tensor output for machine learning where M is a 1D tensor length, U is a 2D image tensor height, V is 2D image tensor width, and N is the number of layers.

Data Product	Tensor Size	Typical Size	Datatype
High pyrometer signal	M x N	40 x 5	PyTorch Float32
Low pyrometer signal	M x N	40 x 5	PyTorch Float32
TED signal	M x N	20 x 5	PyTorch Float32
AE signal	M x N	80 x 5	PyTorch Float32
Coaxial melt-pool images	U x V x N	32 x 24 x 5	PyTorch Float32
After-spread off-axis images	U x V x N	32 x 32 x 5	PyTorch Float32
After-scan off-axis images	U x V x N	32 x 32 x 5	PyTorch Float32
Photometric-stereo heightmaps	U x V x N	32 x 32 x 5	PyTorch Float32
Pore label	1	1	PyTorch Binary

The data products described in this chapter form the experimental basis for the remainder of the dissertation. Chapter 5 describes coordinate mapping, part segmentation, sensor synchronization, XCT segmentation, pore registration, and localized label generation. Chapter 6 reports model results for the synchronized and labeled dataset, and Chapter 7 discusses the implications and limitations of those results.

4.12 Chapter Summary

This chapter described the LPBF experimental system used to generate the multimodal dataset for localized XCT-detected porosity prediction. Nominally dense 316L stainless steel coupons were printed on an Aconity Mini using m4p 316L.03 powder, argon shielding gas, oxygen concentration below 300 ppm, and a 316L stainless steel build plate. Six CLI geometries were reused across two process-parameter groups to produce twelve coupons, with each programmed geometry containing both contour and hatch vectors. The multimodal acquisition system included PrintRite3D process and pyrometer data, acoustic emission, laser-trigger capture, coaxial melt-pool imaging, off-axis imaging, and photometric-stereo heightmap generation. SMMVpy2 coordinated layer tracking, printer pauses, image capture, heightmap processing, and structured data storage. These data provide the experimental basis for process-coordinate synchronization, XCT registration, localized pore labeling, feature extraction, and PMMSF model training.

Chapter 5 Data Preparation

5.1 Overview

This chapter describes the data-preparation workflow used to convert the multimodal LPBF acquisition outputs and post-process XCT measurements into a synchronized, registered, and label-ready dataset for localized pore prediction. Chapter 4 described the LPBF platform, process matrix, sensor suite, acquisition software, and raw or derived data products generated during printing. The present chapter begins with those acquisition outputs and describes the post-processing steps required to place the process data, sensor data, image products, XCT pore morphology, and localized learning labels into a common process-coordinate framework.

The raw acquisition outputs were not directly usable for localized pore labeling because they were stored in different coordinate frames, sampling structures, and time bases. The PrintRite3D/Polaris process files contained laser-position data, process-state signals, and pyrometer channels. The NI myDAQ stored acoustic-emission and laser-trigger signals on a separate acquisition time base. The coaxial camera stored layer-wise image chunks. The off-axis Basler camera stored directional image products in pixel coordinates. The CLI files stored the programmed scan-vector structure. In addition, XCT data were reconstructed in a separate voxel-coordinate frame and had to be segmented, scaled to physical units, and registered to the process-coordinate system before they could be used as pore labels.

The first part of this chapter describes process-coordinate preparation. This includes machine-coordinate calibration, laser-state and coordinate delay correction, CLI scan-context handling, part segmentation from high Laser Drive points, AE synchronization using the laser trigger, coaxial image synchronization, off-axis image registration, photometric-stereo product organization, and synchronized HDF5 dataset construction. These steps convert the raw acquisition products into a common process-coordinate representation that can be queried by part, layer, and location. The overview of registration is shown in Figure 5.1.

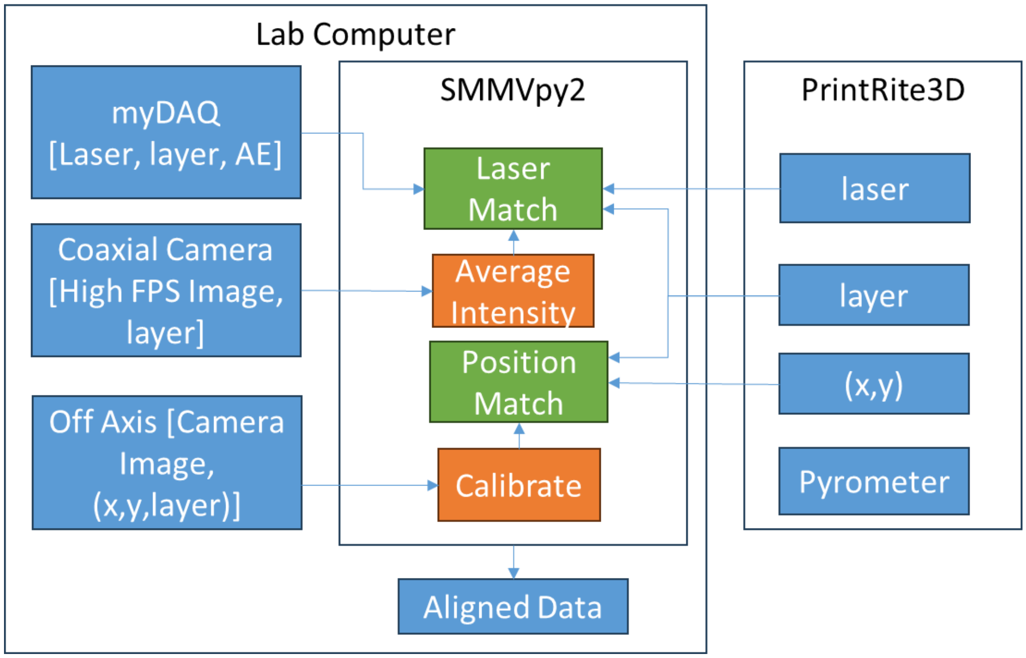

Figure 5.1. Data collected on the lab computer is aligned with the PrintRite3D system. The AE sensor is synchronized by matching laser trigger and layer number with laser trigger collected on the myDAQ. The coaxial camera is matched by taking the average intensity of the image and matching it with the laser trigger. The off-axis camera is aligned with the collected laser coordinate via the camera calibration and layer number.

The second part of this chapter describes the XCT-to-label workflow. XCT scan volumes were separated into individual coupon volumes, exported as image stacks, segmented into pores, filtered into accepted three-dimensional pore components, converted from voxel indices to physical coordinates, registered to the process-coordinate frame, and converted into localized query-volume labels. This workflow produces the registered pore morphology and binary pore/no-pore labels used for model training and evaluation in later chapters.

For consistency with the processing files and existing HDF5 field names, the mapped machine coordinates (build-plate coordinates) fields are referred to as “X-Machine” and “Y-Machine”. In this dissertation, these machine-coordinate fields correspond to the build-plate/process coordinates where the laser reaches the powder bed. Raw laser-position fields refer to the scanner-position data before coordinate mapping. The output of this chapter is therefore an HDF5-based dataset containing mapped process coordinates, part-boundary metadata, synchronized sensor associations, registered image products, registered XCT pore morphology, and localized pore labels.

5.2 Data Streams and Required Post-Processing Outputs

The post-processing workflow was organized around four required outputs: mapped process coordinates, coupon boundaries, sensor-to-process synchronization, and structured HDF5 storage. Chapter 4 documented the acquisition hardware, acquisition settings, process matrix, and raw data products. This data-preparation chapter therefore focuses on the transformations needed to convert those acquisition products into a synchronized process-coordinate dataset.

The main post-processing inputs were the PrintRite3D/Polaris process data, the Laser Drive signal, the CLI files, the NI myDAQ AE and trigger data, the coaxial image chunks, the off-axis Basler image products, and the HDF5 files used to store synchronized outputs. These streams were processed to generate mapped machine-coordinate fields, coupon boundary metadata, trigger-defined AE windows, process-to-image associations, registered off-axis image products, and synchronization metadata.

The physical interpretation of each sensor was reviewed in Chapter 2, and the acquisition configuration was described Chapter 4. The following sections therefore focus on coordinate mapping, part segmentation, synchronization logic, and HDF5 organization.

5.3 Machine-Coordinate Calibration and Process Coordinate Fields

The synchronized process dataset uses X-Machine and Y-Machine as the common XY coordinate basis. These fields define machine coordinates (on the build plate) in millimeters and are used for part segmentation, sensor association, off-axis image localization, and later XCT registration. Note that a separate laser coordinate is captured by the PrintRite3D system. The coordinate preparation workflow therefore required both (machine and laser) process-coordinate fields in the HDF5 data and an image-to-machine calibration for the off-axis Basler camera.

The off-axis camera is mounted to the lid of the LPBF machine. Because opening and closing the lid can slightly shift the camera pose relative to the build plate, the off-axis image calibration was performed before each print. This per-print calibration ensured that Basler image coordinates were mapped to the machine-coordinate frame used by the process data.

The off-axis calibration workflow was implemented using a python script. The script defines a two-dimensional total-degree polynomial model, Poly2D, for mapping paired coordinate points. Two calibrations are performed: Basler image to build plate coordinates and build plate coordinates to laser coordinates.

The image of machine step maps non-distorted calibration coordinates into machine coordinates using a fitted polynomial relationship, as summarized by Equation 5.1 - 5.3.

$$X_m = f_x(x_i, y_i), \qquad Y_m = f_y(x_i, y_i) \tag{5.1}$$

$$X_m = a_0 + a_1 x_i + a_2 y_i + a_3 x_i^2 + a_4 x_i y_i + a_5 y_i^2 \tag{5.2}$$

$$Y_m = b_0 + b_1 x_i + b_2 y_i + b_3 x_i^2 + b_4 x_i y_i + b_5 y_i^2 \tag{5.3}$$

Where x_i, y_i are the Basler image coordinates, X_m, Y_m are the machine coordinates in mm, f_x, f_y are the fitted calibration functions, and $a_0 - a_5$ and $b_0 - b_5$ are the fitted coefficients for the mapping.

The implemented degree-2 calibration produced a two-dimensional machine-coordinate RMSE of 0.00395 mm. This value represents the root-mean-square Euclidean residual between the known machine-coordinate calibration points and the corresponding machine-coordinate locations predicted by the fitted nondistorted-grid-to-machine-coordinate mapping. Since the pixel resolution is about 5µm, the error is deemed low enough. The RMSE definition is given in Equation 5.4.

$$\mathrm{RMSE}_{\mathrm{machine}} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\left(X_{m,i} - \hat{X}_{m,i}\right)^2 + \left(Y_{m,i} - \hat{Y}_{m,i}\right)^2\right)} \tag{5.4}$$

Where N is the number of calibration points, i is the index of the current calibration point, $(X_{m,i}, Y_{m,i})$ is the true location of the ith calibration point, and $(\hat{X}_{m,i}, \hat{Y}_{m,i})$ is the predicted location of the ith calibration point.

After the Basler image to build plate coordinates were mapped, a python script was developed to fit forward and inverse polynomial transforms between Basler image coordinates and machine coordinates. The forward transform maps image pixels into machine coordinates, while the inverse transform supports localization, cropping, and image warping around machine-coordinate positions. The resulting calibration object supports registered off-axis image crops and photometric-stereo products in the same coordinate framework as the process data.

A CLI file was generated with known points to facilitate the calibration was generated with Autodesk Netfabb 2024 shown in Figure 5.2. The file consisted of 121 circles or an 11x11 grid. Each circle had a diameter of 100µm, and the laser spot size was set to 40µm. An example image of a zoomed in calibration point is shown in Figure 5.3. Thresholding is then applied to the image and connected components are used to isolate each point. The mean of each point is then computed and used to represent the center of the point. The point center is then used for calibration.

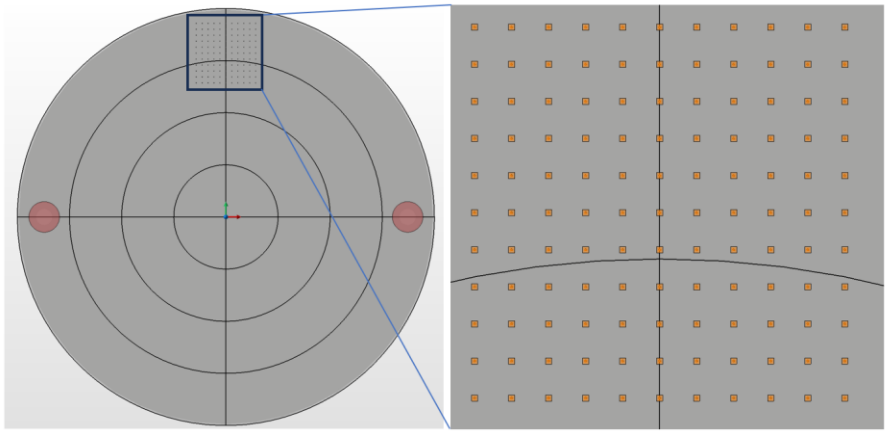

Figure 5.2. Calibration point location on build plate (left) and zoomed in (right)

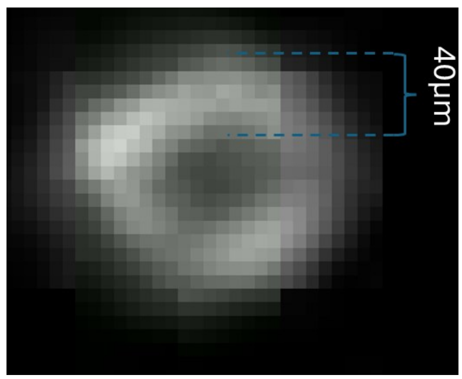

Figure 5.3. Single calibration points for off axis camera.

The grid was designed to encompass as much of the field of view of the camera as possible as shown in Figure 5.4. The locations of the points are defined in Table 5.1.

Table 5.1. Calibration point locations for Basler image calibration

	Min (mm)	Max (mm)	Spacing (mm)
x	-10	10	2
y	45	65	2

The calibration image used for this per-print mapping is shown in Figure 5.4. This image provides the calibration basis for detecting known laser locations and relating the Basler image frame to the machine-coordinate frame. The resulting mapped calibration points are shown in

Figure 5.5. This quality check verifies that the coordinate mapping preserves the expected calibration-point layout and provides a quantitative basis for the reported machine-coordinate RMSE. From this image we can also compute the pixel resolution. The width of the box extents is 3940 pixels, and the physical size is 20mm yielding a 5µm/pixel resolution in the x direction. In the y-direction the extents are 3340 pixels and a height of 20mm yields a 6µm/pixel resolution.

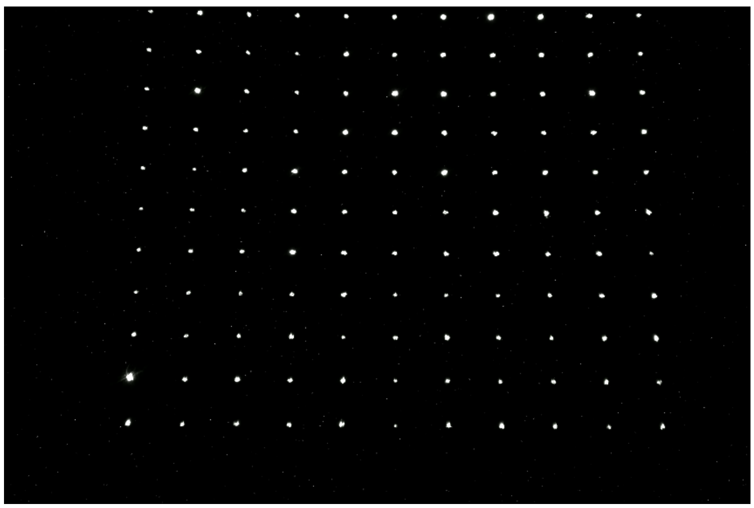

Figure 5.4. Basler calibration image used for machine-coordinate calibration. Known laser locations were imaged using the Basler off-axis camera under extended exposure. Because the camera is mounted to the machine lid, this calibration was performed before each print to account for small camera-position changes introduced when the lid was opened and closed.

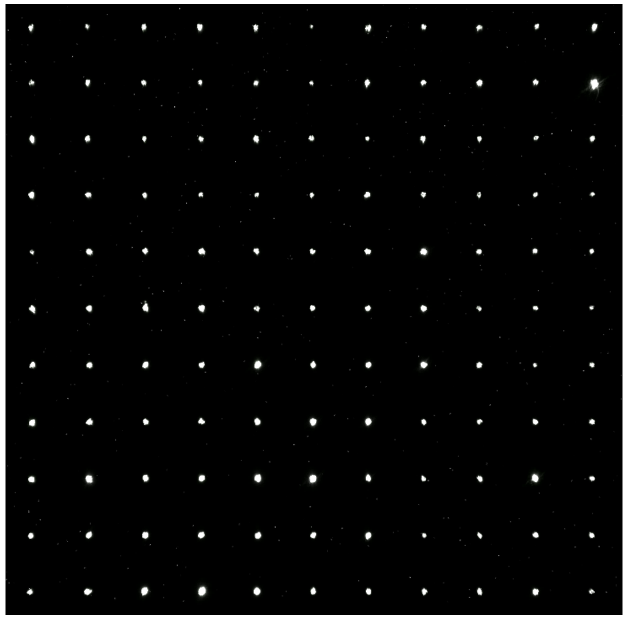

Figure 5.5. Machine-coordinate calibration quality check. Calibration points are mapped into the machine-coordinate frame after applying the nondistorted-grid-to-machine-coordinate mapping. The degree-2 fit produced a two-dimensional machine-coordinate RMSE of 0.00395 mm.

It is noted that the Basler camera is mounted on the machine lid and moving the lid results in camera movement. Therefore, prior to each build the camera must be recalibrated. If the calibration points are outside the bounds of the camera, then one must iterate in calibration and camera movement until the calibration grid is within view. It is also noted that the camera must be focused prior to each build and moving the camera to encompass the grid results in slight defocusing in interest. Therefore, the camera focus must be rechecked after each movement.

5.4 Laser-State and Coordinate Delay Correction

The synchronized LPBF process-monitoring dataset contains laser-state, laser-position, mapped build-plate coordinate, acoustic-emission trigger, and auxiliary sensor channels stored in a common HDF5 structure. Because these channels are used to map process histories to XCT-detected pores, the temporal registration between the laser-state signal and the coordinate channels is critical. A small timing mismatch can place the apparent laser-on location upstream or downstream of the actual melt-track position, which can affect pore-level feature assignment (Yeung and Grantham; Hocker et al.; Kim et al.).

A timing correction was applied using the PrintRite3D delay settings associated with the acquisition system. The manufacturer settings were secToDelayXY = 0.000208 s for the XY coordinate channels and secToDelayLDS = 0.000103 s for the laser-drive or laser-detection signal channels. It is noted that we assume the manufacturer setting is optimal. Errors in this alignment may result in classification errors down the line.

The XY delay was applied to the X-Position, Y-Position, X-Machine, and Y-Machine coordinate rows. The LDS delay was applied to Laser Drive prior to synchronization with the NI laser-trigger channel. The matched edge data index field was also updated after the LDS delay so that stored matched-edge indices remained aligned with the shifted laser-drive edge locations. All other datasets and non-target channels were preserved and unmoved. A visual representation of the delay is shown in Figure 5.6.

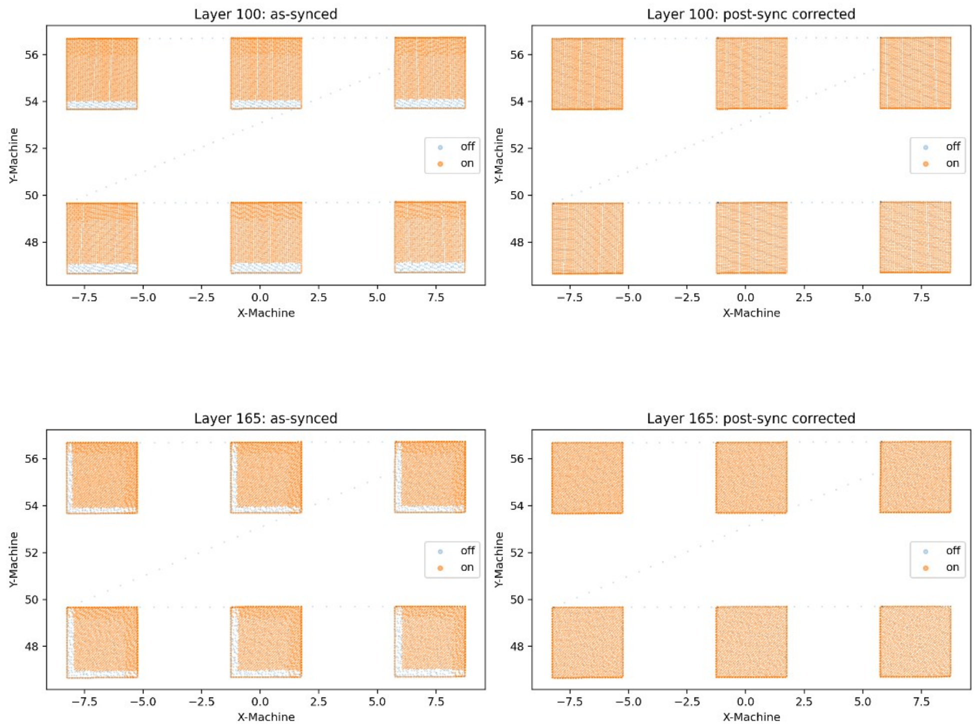

Figure 5.6. Before and after X-Position, Y-Position, and laser drive delays are applied.

Separate XY and LDS delays were applied because the coordinate and laser-state channels do not necessarily represent the same physical instant when they are stored in the synchronized file. The XY channels describe the reported scanner or position trajectory, whereas the LDS and laser-drive channels describe commanded or detected laser-state timing. These signals can have different delay paths because of scanner-control timing, galvanometer response, signal acquisition, interpolation, and synchronization processing.

The correction was evaluated by comparing the laser-on trajectory before and after the delay update in both mapped build-plate coordinates and raw laser-position coordinates. Before correction, laser-drive transitions were visibly offset from coordinate turnarounds and scan-vector endpoints. After applying the PrintRite3D delay settings, the laser-on coordinate paths aligned more closely with the expected scan geometry. The same timing behavior was observed in X-Machine/Y-Machine and X-Position/Y-Position representations, indicating that the mismatch was associated with channel timing rather than only with the coordinate-mapping transform.

This correction should be interpreted as a manufacturer-setting-based timing adjustment to the synchronized process-monitoring data. It does not modify the XCT coordinate frame, the XCT pore locations, or the physical pore morphology. Its purpose is to place the laser-state and coordinate channels on a more consistent time basis before assigning process histories to XCT-detected pore locations. The unmodified synchronized file is preserved for traceability, while the delay-corrected file is used for downstream pore-feature assignment and model training.

The correction does not imply that all residual timing error is removed. Some start/end asymmetry may remain because scanner acceleration, laser on/off dynamics, and interpolation can produce edge-dependent behavior rather than a perfectly uniform delay. However, applying PrintRite3D delay settings provides a physically justified and manufacturer-supported timing correction before assigning laser exposure histories to XCT pore locations.

5.5 CLI Scan and Layer Context

The CLI files provided the programmed layer and scan-vector structure for each coupon. Each CLI file was generated with Autodesk Netfabb Premium 2024. These files contained contour and hatch vectors, layer information, and scan-vector ordering. Chapter 4 described the programmed coupon geometry and process matrix; in this chapter, the CLI files are used only as scan-context information.

The CLI data was not used to impose final part boundaries. Instead, final part segmentation was performed from the collected process data after coordinate mapping. This distinction is important because CLI data describe the intended toolpath, while high Laser Drive points identify where laser activity was recorded in the mapped process data.

The programmed hatch orientations varied across layers, with approximate scan directions of 0 degrees, 45 degrees, 90 degrees, and 135 degrees. These orientations provide context for interpreting layer-wise process data and synchronization behavior. However, contour and hatch points were not separately labeled in the synchronized HDF5 dataset. The synchronized data structure instead preserves process location, layer identity, part boundary, and sensor association.

5.6 Part Segmentation from High Laser Drive Points

After X-Machine and Y-Machine were available in the process data, coupon boundaries were identified from high Laser Drive points. The objective was to define the XY region corresponding to each printed coupon so that process samples and synchronized sensor data could be filtered by part.

Part segmentation was performed using grid-based connected-component clustering of high Laser Drive occupied cells. This method was used instead of k-means clustering. The segmentation workflow reads the mapped HDF5 layer groups, identifies high Laser Drive samples, converts their XY locations into occupied grid cells, groups connected occupied cells, and writes a root-level /part boundaries dataset.

The Laser Drive state was converted into a binary laser-on indicator using Equation 5.5.

$$\delta_i = \begin{cases} 1, & L_i \geq T_L \\ 0, & L_i < T_L \end{cases} \tag{5.5}$$

Each laser-on point was assigned to an occupied XY grid cell using Equation 5.6.

$$c_{x,i} = \text{floor}\left(\frac{X_{m,i} - X_0}{\Delta_g}\right), \qquad c_{y,i} = \text{floor}\left(\frac{Y_{m,i} - Y_0}{\Delta_g}\right) \tag{5.6}$$

Connected occupied cells were grouped using connected-component clustering. Because the coupons are separated in the XY plane, each retained connected component corresponds to one coupon. Conservative rectangular bounds were computed from the occupied grid-cell extents using Equation 5.7-5.10.

$$X_{p,\min} = X_0 + \Delta_g \min_{(c_x,c_y)\in C_p} c_x \tag{5.7}$$

$$X_{p,\max} = X_0 + \Delta_g \left(\max_{(c_x,c_y)\in C_p} c_x + 1 \right) \tag{5.8}$$

$$Y_{p,\min} = Y_0 + \Delta_g \min_{(c_x,c_y)\in C_p} c_y \tag{5.9}$$

$$Y_{p,\max} = Y_0 + \Delta_g \left(\max_{(c_x,c_y)\in C_p} c_y + 1 \right) \tag{5.10}$$

The default conservative grid-cell bounds were used. No manual correction, expansion, or padding was applied. Contour points were included because segmentation used all high Laser Drive occupied cells. Part IDs were assigned consistently with the programmed coupon order.

Algorithm 5.1 summarizes the Part segmentation from high Laser Drive occupied cells

1. Read layer groups from the mapped HDF5 process file.
2. Locate the Laser Drive, X-Machine, and Y-Machine fields.
3. Identify laser-on samples from the Laser Drive state.
4. Convert laser-on process coordinates into occupied XY grid cells.
5. Group occupied cells using connected-component clustering.
6. Remove components below the minimum point-count or cell-count criteria.
7. Compute conservative rectangular bounds for each retained component.
8. Assign part IDs using a consistent spatial order.

9. Write the /part boundaries dataset at the root of the HDF5 file.

10. Save a validation plot showing occupied cells and detected part boxes.

The inspected segmented HDF5 file contained six detected parts. The grid size was 0.05 mm, the bound mode was grid-cell conservative, and the stored Laser Drive threshold was 2.5 V. The generated /part boundaries dataset contains the assigned part ID, conservative X/Y bounds in millimeters, boundary width and height, the number of high Laser Drive points, the number of layers represented, and the number of occupied grid cells. The part-segmentation output is shown in Figure 5.7. This figure verifies that high Laser Drive occupied cells form six separated coupon regions and that the generated conservative rectangular bounds are consistent with the expected print layout.

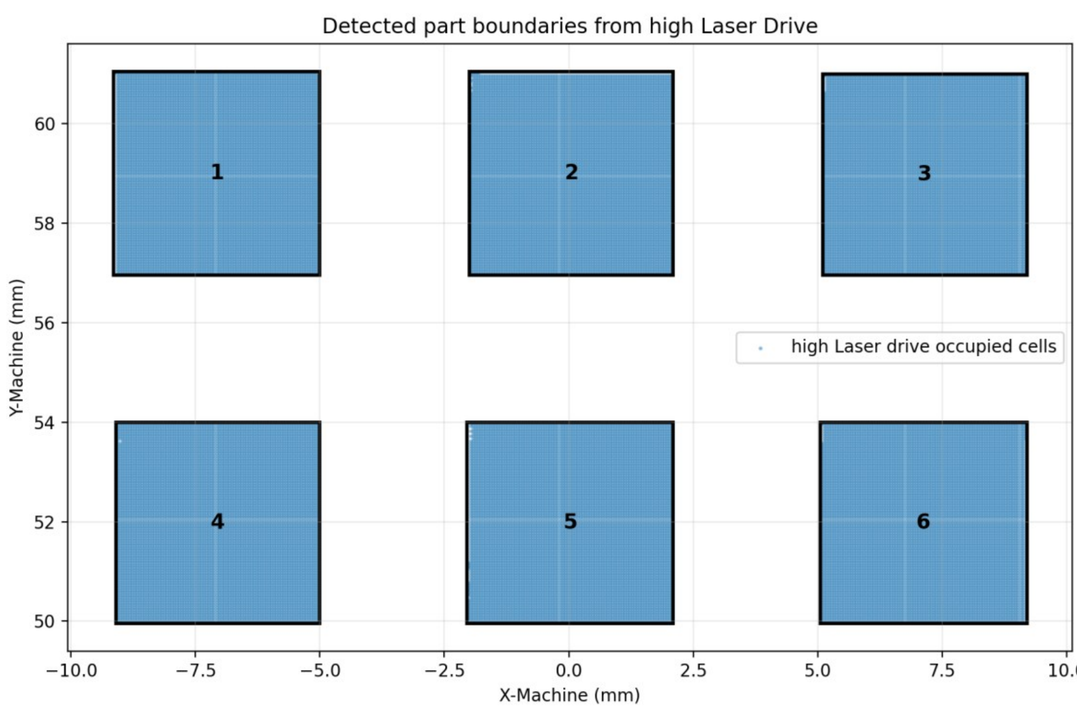

Figure 5.7. Part-segmentation quality check from high Laser Drive points. High Laser Drive occupied cells were identified in the mapped X-Machine/Y-Machine plane and grouped into connected regions corresponding to the six coupons in the print. Conservative rectangular boundaries were generated from occupied grid cells and assigned part IDs consistent with the programmed coupon order.

5.7 AE Synchronization Using the Laser Trigger

AE synchronization was performed using the laser-trigger signal recorded on the same NI myDAQ as the AE signal. This shared acquisition time base allowed acoustic data to be associated with laser-on intervals. The trigger signal was acquired as an analog voltage, with the original 10 V trigger reduced to approximately 1 V using a voltage divider before measurement.

The trigger state was defined using Equation 5.11.

$$\delta_{\mathrm{trig}}(t_i) = \begin{cases} 1, & V_{\mathrm{trig}}(t_i) \geq T_{\mathrm{trig}} \\ 0, & V_{\mathrm{trig}}(t_i) < T_{\mathrm{trig}} \end{cases} \tag{5.11}$$

Consecutive high-state samples define laser-on windows, as shown in Equation 5.12.

$$W_j = \left[t_{\mathrm{on},j}, t_{\mathrm{off},j}\right] \tag{5.12}$$

AE segments were extracted over these laser-on windows. To associate AE samples with process coordinates, laser position was interpolated over the corresponding interval. For an NI sample located between two process-position samples, the interpolated machine coordinate is defined by Equation 5.13.

$$\mathbf{r}_m(t_i) = \mathbf{r}_{m,k} + \frac{t_i - t_k}{t_{k+1} - t_k}\left(\mathbf{r}_{m,k+1} - \mathbf{r}_{m,k}\right) \tag{5.13}$$

This interpolation assumes constant scan speed within the laser-on segment. The output of this step is an association between AE windows, trigger-defined laser activity, layer number, and mapped machine coordinate. The AE synchronization quality check is shown in Figure 5.8. This figure verifies that laser-on intervals can be identified from the trigger channel recorded on the NI myDAQ and used to extract AE windows over corresponding process intervals.

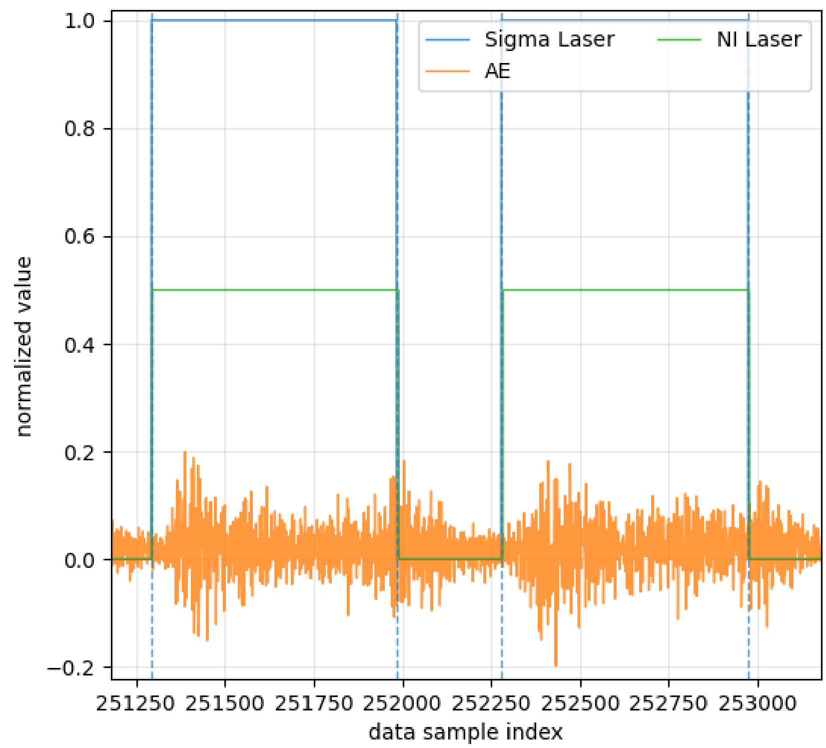

Figure 5.8. AE and laser-trigger synchronization quality check. The laser-trigger channel recorded with AE on the NI myDAQ defines laser-on intervals. AE windows are extracted over trigger-high intervals and associated with interpolated machine coordinates.

5.8 Coaxial Image Synchronization

The coaxial image stream was stored as layer-specific image chunks. Synchronization was performed by matching image-intensity state to laser activity rather than relying on coaxial timestamps as the primary synchronization variable.

For each coaxial frame, the average image intensity was calculated using Equation 5.14.

$$\bar{I}_j = \frac{1}{HW} \sum_{v=1}^{H} \sum_{u=1}^{W} I_j\,(u, v) \tag{5.14}$$

The frame-average intensity sequence was separated into low and high states. The high-intensity state corresponds to laser-on frames, while the low-intensity state corresponds to background or laser-off frames. The image-state sequence was then matched to laser activity in the process data, producing the process-sample-to-image-frame association summarized in Equation 5.15.

$$j_i = g(i) \tag{5.15}$$

When a nearest-frame association was needed, the matched frame was selected using Equation 5.16, subject to consistency between image state and laser state.

$$j_{\text{near}}(i) = \underset{j}{\operatorname{argmin}} \left| t_{\text{process}}(i) - t_{\text{image}}(j) \right| \tag{5.16}$$

The output is a process-sample-to-image association that allows raw coaxial images or image-derived quantities to be retrieved for a mapped process location. The coaxial synchronization quality check is shown in Figure 5.9. This figure verifies that frame-average image intensity can be separated into laser-off and laser-on states and matched to laser activity in the process data.

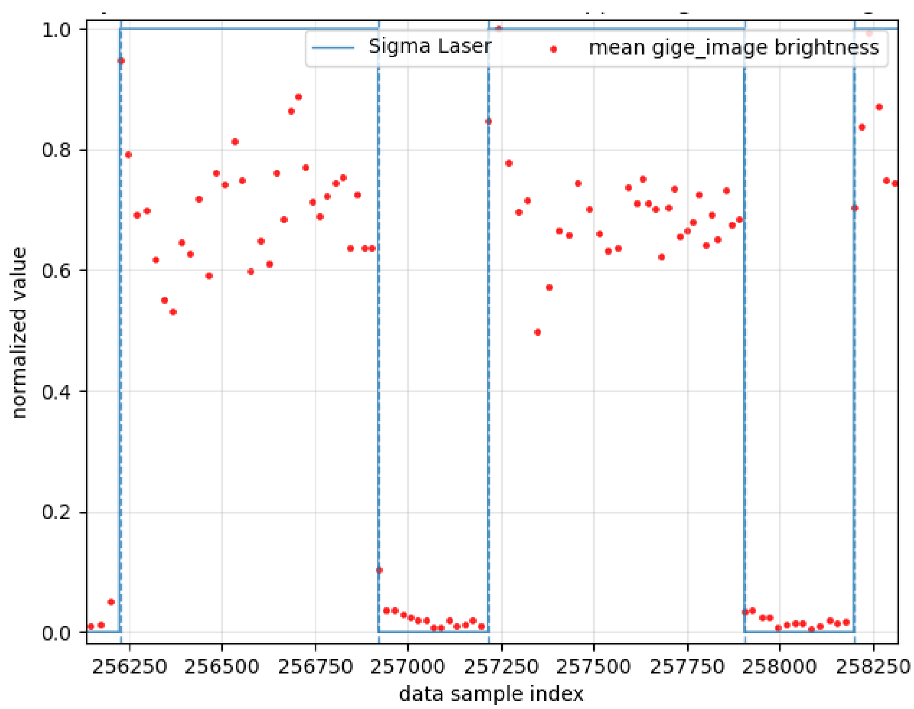

Figure 5.9. Coaxial image synchronization quality check. Coaxial frame-average intensity was separated into low and high states and matched to laser activity in the process data. This matching produced a layer-wise process-sample-to-image association for downstream image retrieval.

5.9 Registered Off-Axis and Photometric-Stereo Products

Off-axis image products were registered to the same machine-coordinate framework used for the process data. This data-preparation chapter described the off-axis acquisition and photometric-stereo workflow; this section describes how those image products are organized after per-print calibration.

The registered off-axis products were cropped or indexed so that they corresponded to the same coupon-coordinate regions used for one-dimensional process data. Because the off-axis camera is mounted to the machine lid, the calibration from Section 5.3 was repeated before each print rather than assumed to be fixed across builds. This calibration produced the machine-coordinate transform used for registered off-axis image crops and photometric-stereo products.

Photometric-stereo heightmaps were treated as derived registered image products, not as separate sensor stream. The relevant output is that the off-axis image and heightmap products are placed into the same machine-coordinate framework as the process data and can be queried using part ID, layer, and XY location. The off-axis registration quality check is shown in Figure 5.10. This figure verifies that off-axis image products can be registered or cropped into the same machine-coordinate frame used for process data and part segmentation.

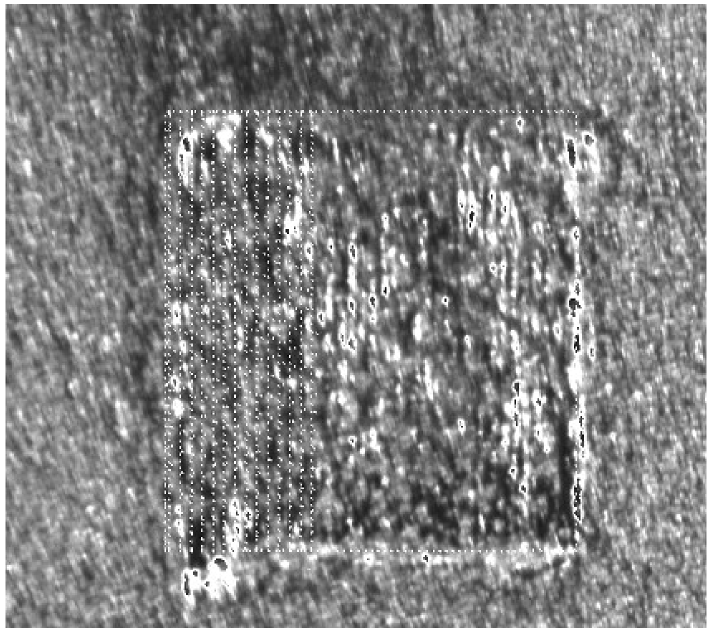

Figure 5.10. Off-axis calibration quality check. Known laser positions were imaged with the Basler camera and used to calibrate the mapping between image coordinates and machine coordinates. The resulting transform supports registered off-axis image crops and photometric-stereo products.

Each image produced was centered on the laser position, cropped to 32x32 pixels and corrected using the inverse transform for Section 5.3. 32x32 pixels correspond to a resampled area of approximately 180μm x 180μm. If the scanned direction is larger than 180μm the image is down sampled. Images were kept to this size due to memory constraints during training.

5.10 Synchronized HDF5 Dataset Organization

The final output is a synchronized HDF5 structure that preserves process coordinates, part boundaries, and sensor associations. Coupon regions are defined by the root-level /part boundaries table. Layer-resolved process data are stored in layer groups. Sensor synchronization outputs are stored as datasets, lookup tables, paths, or references depending on the modality and file size. The layout of the HDF5 file is shown in Table 5.2. An example of the synchronized data is shown in Figure 5.11.

Table 5.2. Synchronized HDF5 outputs generated by the data-preparation workflow.

HDF5 location or output	Description	Purpose
/layer*/data and /layer*/names	Layer-resolved process arrays and corresponding field names	Stores process data and enable field lookup
X-Machine, Y-Machine	Mapped machine-coordinate fields in millimeters	Common XY coordinate basis
Laser drive or Laser Drive	Laser-state/process-state signal	Input for part segmentation and synchronization checks
/part boundaries	Root-level coupon boundary table	Defines part-level XY filtering regions
/part boundaries attributes	Threshold, grid size, bounds mode, source file, and schema metadata	Documents segmentation settings
AE synchronization outputs	Trigger-defined AE windows, features, paths, or references	Associates AE data with laser-on intervals
Coaxial synchronization outputs	Process-to-frame associations or image references	Associates coaxial frames with process samples
Off-axis calibration outputs	Per-print image-to-machine calibration object and registered image products	Supports registered image and heightmap retrieval
Photometric-stereo products	Registered heightmap products or paths	Provides derived topographic image input

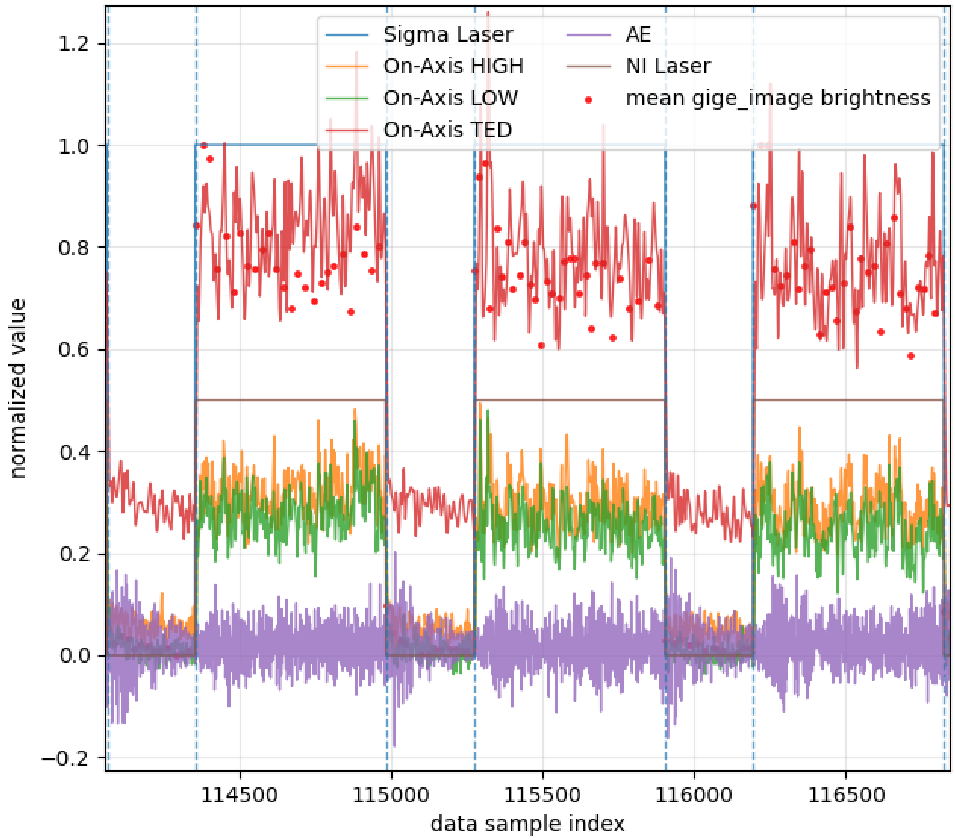

Figure 5.11. Overlay of all collected signal for three scan lines.

5.11 Process-Coordinate Quality Checks

The checks in this chapter verify that coordinate mapping, segmentation, and synchronization were completed consistently before XCT registration. They are presented as quality checks rather than results.

The primary coordinate-calibration check is the machine-coordinate RMSE reported in Section 5.3. The degree-2 nondistorted-grid-to-machine-coordinate calibration produced a two-dimensional RMSE of 0.00395 mm, corresponding to approximately 3.95 µm. This value reports the residual of the calibration fit in machine-coordinate units and does not represent XCT registration error or pore-labeling uncertainty.

The part-boundary validation figure is a central quality check because it confirms that high Laser Drive occupied cells form six separated coupon regions and that the resulting rectangular bounds match the expected print layout. The AE/trigger quality check verifies that laser-on intervals are identifiable in the NI time frame. The coaxial synchronization quality check verifies that image-state transitions align with laser activity. The off-axis calibration quality check verifies that image products can be queried in the machine-coordinate frame.

5.12 XCT-to-Label Workflow Overview

The process-coordinate preparation workflow established the laser-position data, synchronized sensor data, part boundaries, and layer-wise process information needed for XCT-based label generation. The XCT workflow converts post-process pore measurements into localized ground-truth labels in the same coordinate frame as the in-situ process data.

The workflow begins with XCT scan data that have been separated into individual coupon volumes and ends with registered pore morphology and localized pore/no-pore labels for model evaluation in Chapter 6. The major operations are XCT volume preparation, pore segmentation, three-dimensional pore filtering, XCT-to-process registration, registered pore-table consolidation, and localized query-volume label generation. The learning workflow uses these registered labels to construct sensor-aligned datasets for supervised learning and sensor-fusion analysis.

5.12.1 XCT Scanning, Part Isolation, and Volume Export

XCT scanning was performed by Cloud NDT in Irvine, California. They utilized a NSI X5000 CT system, an industrial 3D X-ray/CT inspection system, running at 220kV and exposed the samples for 3 hours. Samples were scanned in sets of four coupons, with each four-coupon scan set requiring approximately three hours. The coupons were held in foam during scanning.

The reconstructed voxel size is treated as approximate and is reported using the value used elsewhere in the workflow, approximately 0.0076mm.

The four-coupon XCT volumes were processed in VGStudio to crop and separate the individual coupon volumes. VGStudio pore analysis was then used to detect the pores in the samples. The working file format was VGL. MyVGL was then used to export image stacks from the separated coupon volumes.

5.12.2 XCT Pore Segmentation and Classified Volume Generation

The exported XCT image stack was processed using VGStudio. The slices were exported from VGStudio, and a Python script was developed to convert the image stack into a NumPy array. The images were loaded sequentially, and thresholding was applied to distinguish pores from the bulk material. A second threshold separated the bulk material from the exterior background. The classified volume was stored as an 8-bit integer NumPy array. A representative segmentation example is shown in Figure 5.12.

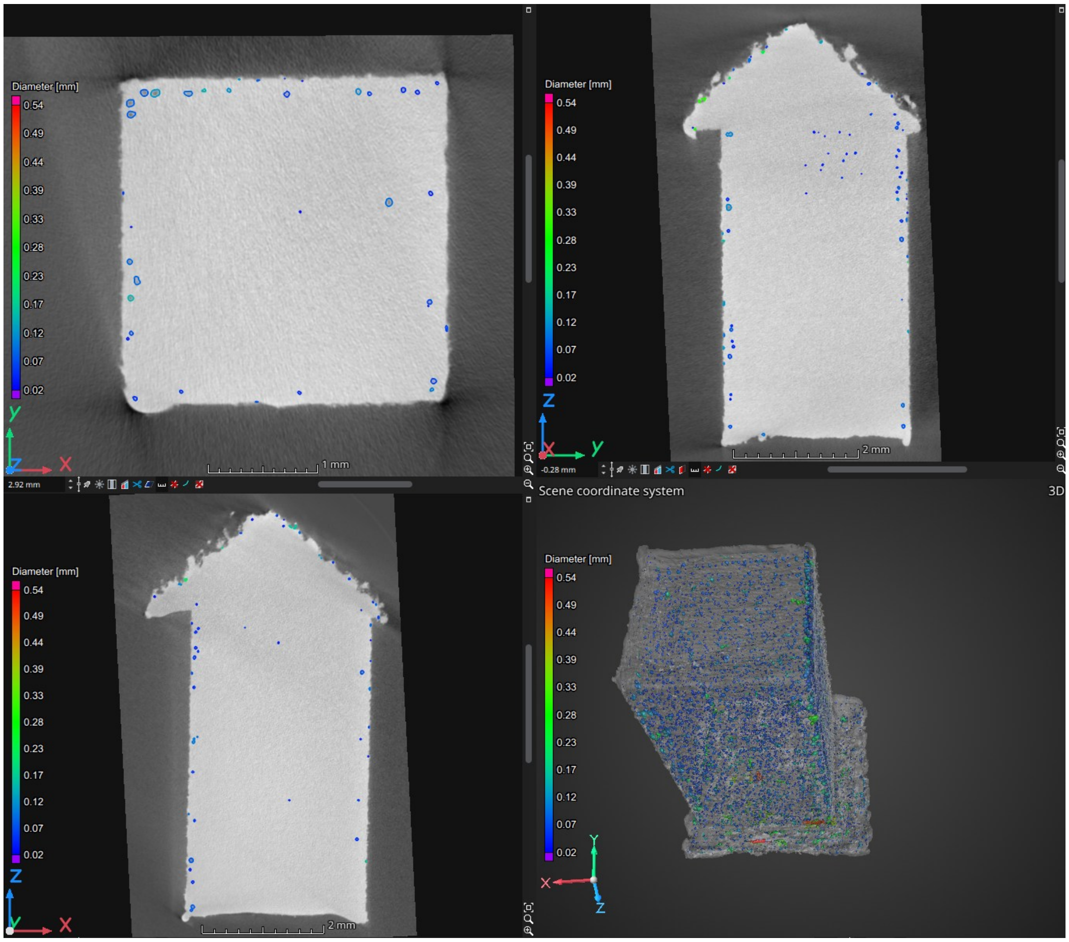

Figure 5.12. VGStudio pore segmentation with pores marked.

The final classified volume used the following integer encoding exterior background = 0, pore voxels = 1, and solid coupon material = 2. This convention is a voxel-class representation and should not be confused with the binary machine-learning labels generated later in the workflow. An example of an XCT layer to registered NumPy array is shown in Figure 5.13.

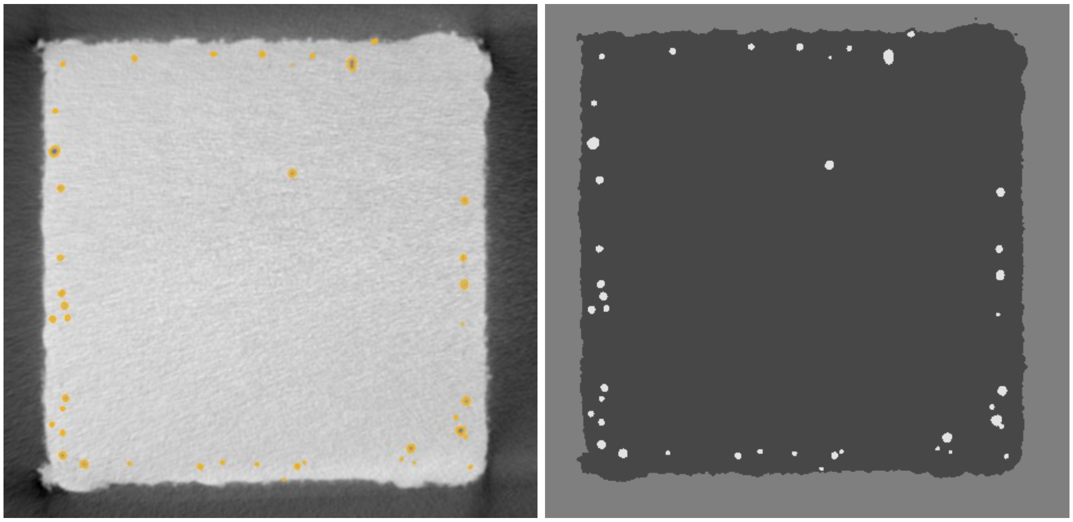

Figure 5.13. Pore selection (right) with pore exterior marked in yellow and saved NumPy layer (left) with pores marked in white, light gray background, and dark gray condensed part.

5.12.3 XCT Voxel Coordinates and Local Physical Coordinates

The separated XCT image stack defines a local voxel coordinate system. Voxel indices must be converted to physical coordinates before the pore field can be registered to the process-coordinate frame. The voxel-to-physical conversion uses the reconstructed voxel size and maps slice and in-plane pixel indices to local XCT coordinates.

The voxel-to-local-coordinate conversion is represented by Equation 5.17. Pore volume is computed from accepted voxel count and voxel volume using Equation 5.18, and equivalent spherical diameter is computed using Equation 5.19. These quantities provide the physical pore-size measures stored in the registered pore table.

Let a voxel center be indexed by slice and pixel coordinates (x_i, y_i, z_i), with voxel spacing (s_z, s_y, s_x) in millimeters. If the local XCT origin is placed at (x_0, y_0, z_0), then

$$\mathbf{x}_{\mathrm{XCT},i} = \begin{bmatrix} X_{\mathrm{XCT},i} \\ Y_{\mathrm{XCT},i} \\ Z_{\mathrm{XCT},i} \end{bmatrix} = \begin{bmatrix} (x_i - x_0)s_x \\ (y_i - y_0)s_y \\ (z_i - z_0)s_z \end{bmatrix} \tag{5.17}$$

Where $\mathbf{x}_{\mathrm{XCT},i}$ is the local physical XCT coordinate of voxel i; s_x, s_y, and s_z are voxel spacings; and (x_0, y_0, z_0) defines the local XCT origin.

For pore component p, with N_p accepted voxels,

$$V_p = N_p s_x s_y s_z. \tag{5.18}$$

Where V_p is pore volume and N_p is the accepted voxel count for pore p.

$$d_{\mathrm{eq},p} = \left(\frac{6V_p}{\pi}\right)^{1/3}. \tag{5.19}$$

Where $d_{\mathrm{eq},p}$ is the equivalent diameter of a sphere with volume V_p.

The local XCT coordinate frame is distinct from the process-coordinate frame. Voxel scaling converts XCT indices into local millimeter coordinates, while registration maps those local coordinates into the build/process coordinate system.

5.13 XCT-to-Process Registration

After pore filtering, accepted pore geometry was registered to the process-coordinate frame. The registration workflow maps XCT-local coordinates into machine/process coordinates using 4 x 4 homogeneous transformations as shown in Figure 5.14. The process-coordinate data provides X-Machine and Y-Machine in millimeters, and process Z is computed from the layer number and the layer height. In this dissertation, the layer height is 0.03 mm. The process-coordinate data is filtered using the laser-on data to obtain the coordinates where the laser melted the powder and created the solid part. With XCT coordinates in mm and laser-on coordinates in mm, the data can be aligned.

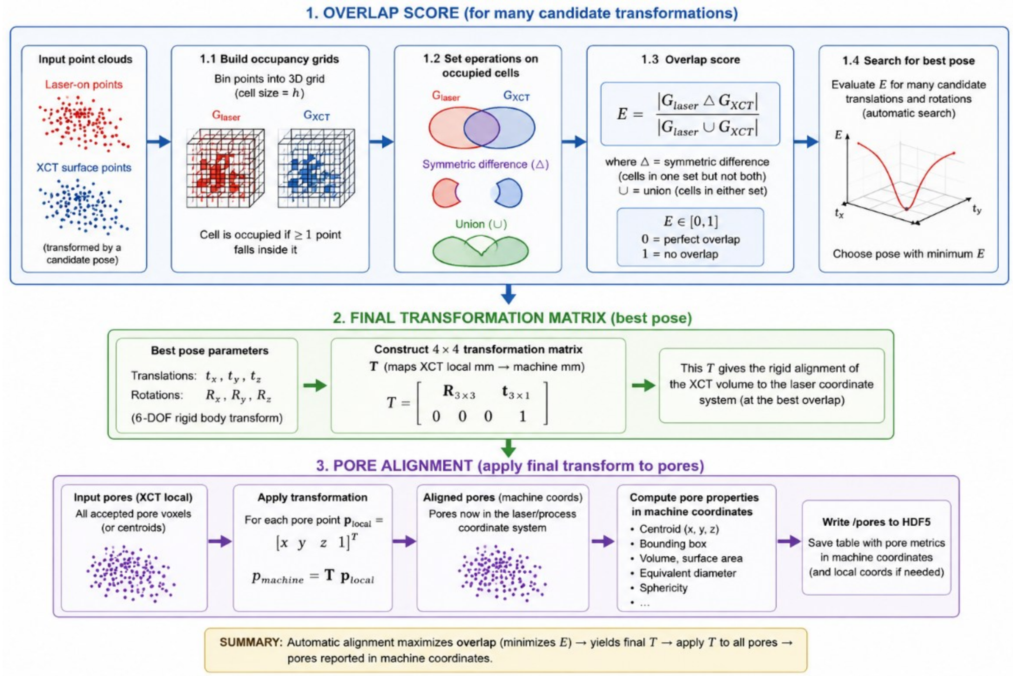

Figure 5.14. Pore alignment process overview

The XCT-to-process transformation is represented by Equation 5.20. This equation defines how a pore location measured in the XCT volume is converted into the same coordinate frame as the synchronized process data.

$$\begin{bmatrix} X_{\text{proc}} \\ Y_{\text{proc}} \\ Z_{\text{proc}} \\ 1 \end{bmatrix} = \mathbf{T}_{\text{XCT}\to\text{proc}} \begin{bmatrix} X_{\text{XCT}} \\ Y_{\text{XCT}} \\ Z_{\text{XCT}} \\ 1 \end{bmatrix}. \tag{5.20}$$

Where $\mathbf{T}_{\text{XCT}\to\text{proc}}$ is the saved 4×4 homogeneous transformation matrix mapping local XCT coordinates into the process-coordinate frame.

$$\mathbf{T}_{\text{XCT}\to\text{proc}} = \begin{bmatrix} R & t \\ 0^T & 1 \end{bmatrix} \tag{5.21}$$

where R is the fitted rotation matrix and t is the fitted translation vector.

To obtain the transformation matrix, a point cloud of the laser-on process points and the XCT points were used to generate a surface grid. 75,000 surface points were used for each. Surface points were used instead of points sampled throughout the entire XCT volume because the surface contains the edges and shape needed to determine the coupon's position and orientation. Surface points are simply those where a voxel shares a face with the background. Interior points are not used because overlapping within the coupon interior could produce a low overall error even when the coupon boundaries remain misaligned.

To perform alignment a registration objective is defined. The registration objective minimized the normalized mismatch between occupied process-grid cells and occupied XCT-grid cells:

$$\mathrm{E} = \frac{|\mathcal{G}_{\mathrm{laser}} \triangle \mathcal{G}_{XCT}|}{|\mathcal{G}_{\mathrm{laser}} \cup \mathcal{G}_{XCT}|} \tag{5.22}$$

where $\mathcal{G}_{\mathrm{laser}}$is the set of occupied grid cells from the laser-on process points, $\mathcal{G}_{\mathrm{XCT}}$is the set of occupied grid cells from the transformed XCT surface points, and $\triangle$ denotes the symmetric difference. The symmetric difference is the set of grid cells occupied by one dataset but not by both. Lower values indicate better overlap between the process-coordinate point cloud and the transformed XCT geometry.

The XCT and laser points are initially aligned by placing the XCT center of mass at the laser-data center of mass. Then using the objective in Equation 5.22, initial orientation candidates were tested, including no flip, 90-degree Z-axis rotations, and 180-degree local X- or Y-axis flips. The best candidate is then fine-tuned with six-degree-of-freedom (6-DOF) alignment using rotations up to 10 degrees with a coarse step of 2 degrees and a fine step of 0.02 degrees, and translations up to 350μm with a coarse step of 10μm and fine step of 5μm. A vertical synchronization step was also applied to match the top of the transformed XCT volume to the top process-layer height. This Z-only correction was computed as:

$$\Delta \mathrm{z} = \mathrm{z}_{\mathrm{top,laser}} - z_{top,xct} \tag{5.23}$$

where $z_{\mathrm{top,laser}}$is the maximum process-layer height for the selected part and $z_{\mathrm{top,XCT}}$is the top-surface reference from the transformed XCT point cloud. The final transformation could then be reviewed and adjusted using the manual WebGL alignment interface as shown in Figure 9.2 in the Appendix.

The vertical synchronization is required due to surface distortions on the top of the part. Specifically, distortions where the laser starts and ends build up every layer until a size around 200μm is achieved as shown in Figure 5.15.

After the final transformation was selected, pore voxels were transformed into the process-coordinate frame. For each accepted pore component r, the registered pore centroid was computed and mapped by:

$$p_{r,process} = \frac{1}{N_r} \sum_{q=1}^{N_r} \mathbf{T}_{\mathrm{XCT}\rightarrow\mathrm{proc}} \begin{bmatrix} x_{r,q} \\ 1 \end{bmatrix} \tag{5.24}$$

where N_r is the number of accepted voxels in pore r, and $\mathbf{x}_{r,q}$is the local XCT coordinate of voxel q in that pore. The workflow also computed pore volume, equivalent diameter, axis-aligned extents, and minimum rotated XY bounding-box fields from transformed voxel-corner geometry. The final /pores table therefore provided registered pore morphology in the same process-coordinate frame used for sensor synchronization, process-parameter assignment, and localized pore-label generation.

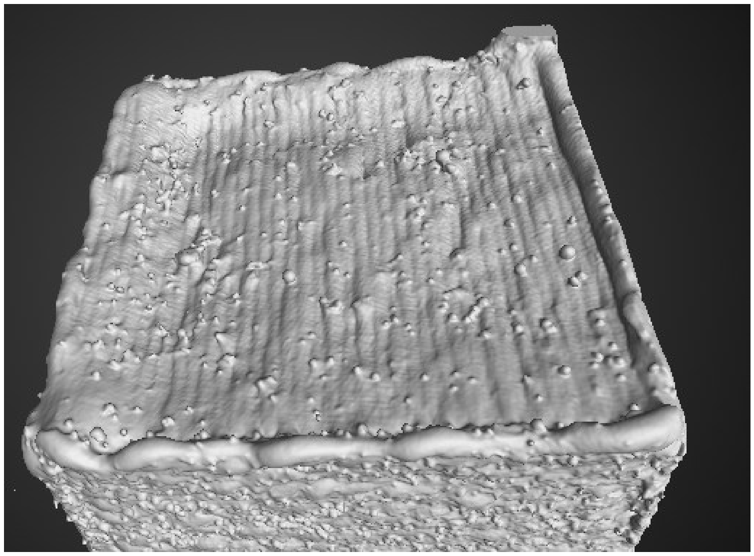

Figure 5.15. Part top edge distortion is observed, so alignment is forced to center of part top

The registered XCT-to-process overlay is shown in Figure 5.16. This figure is the primary geometric quality check for registration because it shows whether the transformed XCT geometry aligns with the process-coordinate part region generated in the process-coordinate workflow.

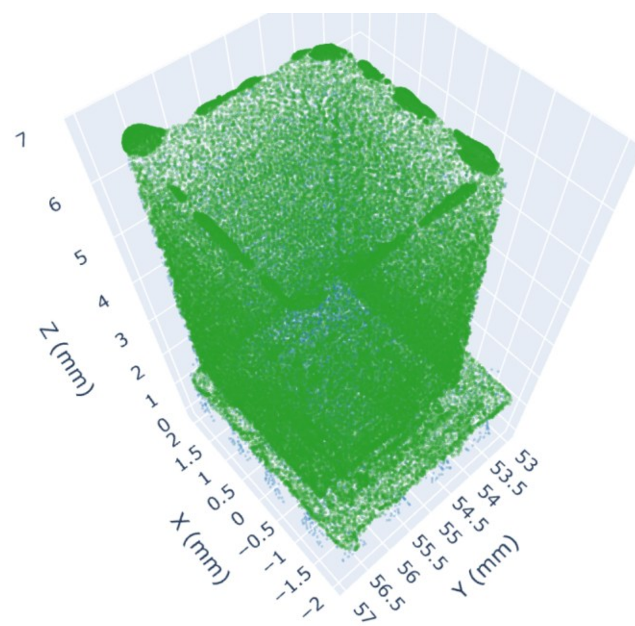

Figure 5.16. XCT-to-process registration overlay. The transformed XCT geometry and pore locations are shown in the process-coordinate frame relative to the laser/process point cloud or part-boundary reference.

Registration quality was evaluated geometrically. Registered pore centroids and pore footprints should fall within the expected part region and layer range. Large systematic offsets in X, Y, or Z would indicate registration error and would directly affect localized pore labels.

5.14 Registered Pore Morphology

The registration workflow produces a registered pore morphology table for each processed part. The table includes pore identity, registered centroid, voxel count, volume, surface area, equivalent diameter, sphericity, axis-aligned bounding-box dimensions, local XCT centroid, and minimum rotated XY bounding-box geometry. The rotated XY bounding-box fields were included to provide a more specific pore footprint for downstream pore-morphology analysis and potential alternative overlap-based labeling.

The registered pore fields used for downstream label generation are summarized in Table 5.3.

Table 5.3. Registered pore morphology fields used for localized label generation.

Field or field group	Description
part_number	Coupon identifier added during multi-part consolidation
pore_id	Pore identifier assigned within the registered part table
centroid_x_mm, centroid_y_mm, centroid_z_mm	Registered pore centroid in process coordinates
voxel_count	Number of accepted XCT voxels in the pore component
volume_mm3	Pore volume computed from voxel count and voxel size
surface_area_mm2	Estimated pore surface area
equiv_diameter_mm	Diameter of a sphere with equivalent pore volume
sphericity	Shape metric based on pore volume and surface area
bbox_x_min_mm through bbox_z_max_mm	Axis-aligned registered pore extents
length_x_mm, length_y_mm, length_z_mm	Axis-aligned pore dimensions
local_centroid_x_mm, local_centroid_y_mm, local_centroid_z_mm	Pore centroid before process-coordinate transformation
min_bbox_center_x_mm, min_bbox_center_y_mm	Center of the minimum rotated XY pore footprint
min_bbox_major_length_mm, min_bbox_minor_length_mm	Major and minor dimensions of the rotated XY pore footprint
min_bbox_angle_deg	Orientation of the rotated XY pore footprint
min_bbox_x0_mm through min_bbox_y3_mm	Corner coordinates of the rotated XY pore footprint
min_bbox_valid	Indicator for whether the rotated XY bounding box is available

5.15 Localized Query-Volume Label Generation

After registration, each accepted pore is represented in the same coordinate frame as the synchronized process data. Localized labels are generated by placing a process-centered query volume at selected laser-on process locations and testing whether registered pore geometry overlaps that volume. Volumes were generated based on a user-provided query volume that consisted of width (transverse to laser path), length (parallel to laser path), height (# of layers), and stride (distance between points on laser path). The laser path is first broken up based on the stride to obtain anchor points. Volumes are built around these points based on the user input. Finally, the volume bounds were used to select the signal data associated with each point from the aligned x- and y-coordinate data. The selected data were then stored in a database for model training. A high-level overview is shown in Figure 5.17.

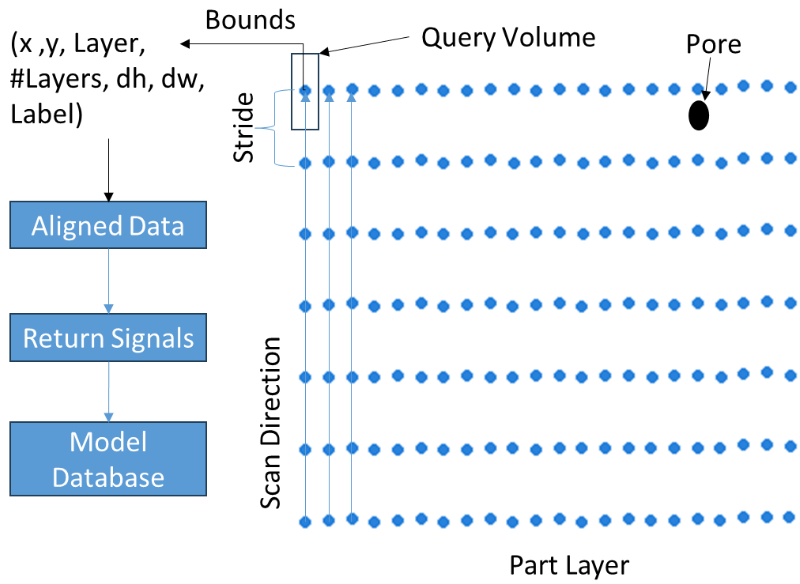

Figure 5.17. Overview of query volume and data generation.

If the width is greater than the laser spot size, a parallel scan line is searched for in the perpendicular direction of the scan direction. One perpendicular direction was searched first. If no candidate scan line was found within one laser-spot diameter, the opposite direction was searched. An overview of this method is shown in Figure 5.18. Note that this method is only tested on simple hatching. For other hatching methods it may not be applicable.

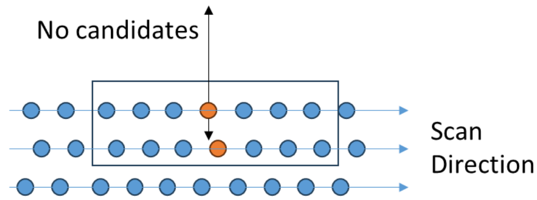

Figure 5.18. Multiple-scan-line query technique. Red represents the anchor points for each row.

The localized labeling operation is illustrated in Figure 5.19. The query volume is centered at the process location in X and Y. Its Z extent is defined using the current layer and the selected number of pore-label layers. With the default current-window-only rule, the same localized XY footprint is projected downward through the current layer and the configured previous label layers.

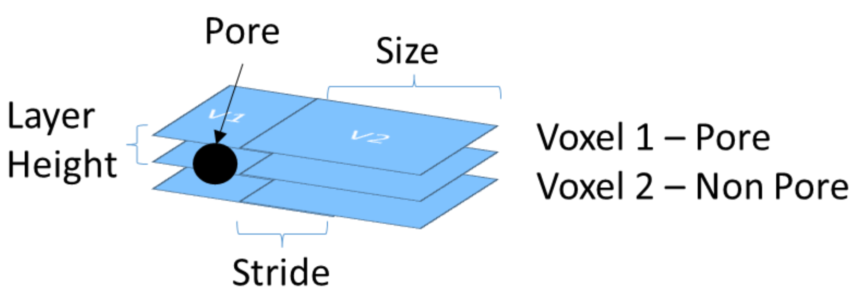

Figure 5.19. Localized query-volume labeling schematic. A query volume is centered on a selected process location. Registered pore geometry is tested against the query volume, and a positive label is assigned when the pore-center rule is satisfied.

The query-volume bounds are defined by Equation 5.25. The query volume of 160 μm transverse × 300 μm scan direction × 150 μm build direction is large enough to encompass most of the pores as shown in Figure 5.20. Larger pores reported are due to the structural material at the bottom of the part which is not used in the analysis. Since the pore size is less than or equal to the query volume, a simple pore center check is performed to determine if a pore is contained within a volume.

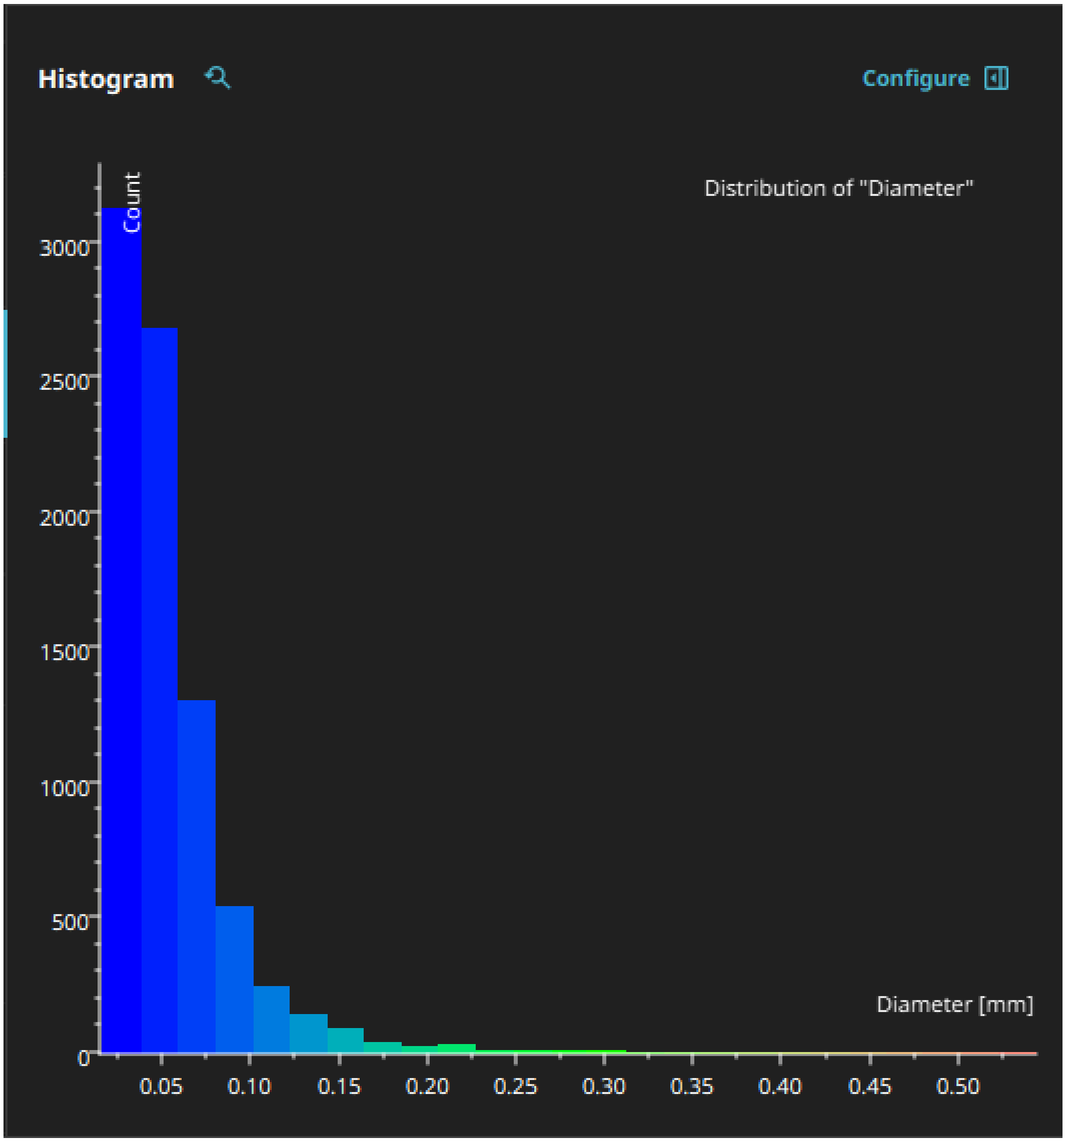

Figure 5.20. Histogram of pore count vs diameter for a representative part.

The software supports both pore-volume-overlap and pore-center labeling criteria. For the fixed query-volume configuration used in the reported results, a positive label was assigned when the registered pore center fell within the query-volume bounds defined by Equation 5.25. The resulting binary rule is given by Equation 5.26.

For a process sample centered at (X_c, Y_c) on layer ℓ, with X window width w, Y window length t, layer height h_ℓ, and pore-label depth of n_p layers,

$$X_{\min} = X_c - \frac{w}{2}, \qquad X_{\max} = X_c + \frac{w}{2},$$

$$Y_{\min} = Y_c - \frac{t}{2}, \qquad Y_{\max} = Y_c + \frac{t}{2},$$

$$Z_{\max} = \ell h_\ell, \qquad Z_{\min} = Z_{\max} - n_p h_\ell. \tag{5.25}$$

Where n_p is the number of pore-label layers included in the query depth.

For pore p with coordinates X_{p}, Y_{p}, Z_{p} and pore label y

$$y = \begin{cases} 1, & (X_{\max} > X_{\mathrm{p}} > X_{\min})(Y_{\max} > Y_{\mathrm{p}} > Y_{\min})(Z_{\max} > Z_{\mathrm{p}} > Z_{\min}) \\ 0, & \text{otherwise.} \end{cases} \tag{5.26}$$

The binary learning label uses the same convention as the classified XCT volume. In the classified XCT volume, pore voxels are encoded as 1. In the learning dataset, a process sample is assigned label = 1 when the query volume satisfies the configured center-based labeling rule and label = 0 otherwise.

5.16 Chapter Summary

This chapter also described the laser-state and coordinate delay correction applied before downstream pore-feature assignment. The correction used PrintRite3D manufacturer delay settings, applied separate XY and laser-drive/LDS timing shifts to the synchronized HDF5 fields, and preserved the original synchronized file for traceability.

This chapter described the conversion of XCT scan data into registered pore labels for localized LPBF pore prediction. XCT scanning was performed by Cloud NDT in Irvine, California using an NSI X5000 CT system. Samples were scanned in four-coupon sets, held in foam, and scanned for approximately three hours per set. VGStudio was used to crop, separate, and export individual coupon image stacks, and VGStudio pore analysis was used as the final pore-analysis method.

VGStudio was used to segment the pores and export the classified coupon image stacks. A custom Python workflow converted the exported image stacks into three-dimensional NumPy arrays for subsequent registration and analysis. The final classified volume encoded exterior background as 0, pore voxels as 1, and solid coupon material as 2.

Accepted pore geometry was registered to the process-coordinate frame using the XCT-to-process synchronization workflow. The resulting registered pore table included centroids, pore volumes, equivalent diameters, axis-aligned extents, and rotated XY pore footprints. Part-specific pore outputs were consolidated into /pores_all.

Finally, registered pore geometry was converted into localized pore/no-pore labels using process-centered query volumes. The registered pore morphology and localized labels generated in this chapter provide the ground-truth basis for the model results reported in Chapter 6.

Chapter 6 Results

6.1 Overview

This chapter presents the quantitative evaluation of the proposed SFML framework for localized XCT-detected pore prediction in LPBF. The results are organized into seven sections: (1) dataset and label statistics, (2) individual sensor performance, (3) multimodal sensor fusion, (4) physics-informed learning, (5) process-informed learning, (6) fusion architecture comparison, and (7) inference and data alignment time.

Unless otherwise stated, all experiments presented in this chapter employ the process-informed configuration introduced in Section 3 to provide a consistent basis for comparing sensing modalities, multimodal sensor fusion, physics-informed learning, and fusion architectures. The contribution of process-informed learning is evaluated separately in Section 6.6 through an ablation study by comparing the proposed framework with and without process-informed weighting.

Because localized pore prediction is an imbalanced binary classification problem, multiple complementary evaluation metrics were used to assess model performance. Validation loss was used to monitor training convergence and select the final model. Model performance was evaluated using accuracy, AUC, F1 score, precision, and recall. Accuracy provides an overall measure of classification performance, whereas the AUC evaluates the model's ability to distinguish pore-containing and non-pore regions independently of the classification threshold. The F1 score was considered the primary evaluation metric because it balances precision and recall, providing a robust measure of pore-detection performance for imbalanced datasets. Precision measures the reliability of predicted pore-containing regions, whereas recall quantifies the fraction of actual pore-containing regions correctly identified.

Validation was performed using a layer-level split within the training data, whereas final performance was evaluated using a held-out coupon (Print 1, Part 4) that was excluded from model development. Consequently, the validation metrics represent performance within the training-domain population, whereas the test metrics evaluate generalization to an unseen coupon fabricated under the same material system, machine, geometry family, and experimental campaign.

Table 6.1 summarizes the final dataset and evaluation protocol used throughout this chapter. The dataset comprises twelve 316L stainless steel LPBF coupons containing 39,114 accepted XCT-detected pore components. Localized pore labels were generated using a pore-center containment criterion within a query volume of 160 μm × 300 μm × 150 μm, establishing the ground truth for supervised learning. Model development was performed using eleven coupons, with validation based on a layer-level split within the training data. Final performance was evaluated using a held-out coupon (Print 1, Part 4) that was excluded from model development. Consequently, the validation metrics represent performance within the training-domain population, whereas the test metrics evaluate the model's generalization to an unseen coupon fabricated using the same material system, machine, geometry family, and experimental campaign. The held-out test set contains 3,305 positive samples and 12,352 negative samples, reflecting the constructed query-sample distribution after the sampling procedure.

Table 6.1. Overview of training, testing, and validation dataset parameters

Category	Value
Prints	2
Coupons	12
Material	316L stainless steel
Query volume (transverse direction, scan direction, build direction)	160 µm × 300 µm ×150 µm
Scan stride	300 µm
Layer Stride	2 Layers
Training & Validation parts	11
Training / Validation split	70% / 30%
Held-out test part	Print 1, Part 4
Split Method	Contiguous Layers
Test positives	3305
Test negatives	12352

All work was performed on a workstation running Windows 11 Enterprise Version 10.0.26200, Build 26200. This system includes two 10 core Intel Xeon Silver 4210R CPUs running at 2.4GHz. 296GB of DDR4 memory was installed on the workstation. Three Nvidia A5000 GPUs are used for training and inference. A Broadcom NetXtreme Gigabit Ethernet adaptor was installed for camera communication.

6.2 Dataset and Label Statistics

Table 6.2 summarizes the number of volume queries containing pores and the measured porosity for each LPBF coupon. Across the twelve 316L stainless steel coupons, the workflow identified a total of 39,114 pore components. Print 1 contained 18,510 samples containing pores (mean: 3,085 pores per coupon), whereas Print 2 contained 20,604 samples containing pores (mean: 3,434 pores per coupon). The samples containing pores per coupon ranged from 2,957 to 3,846 per coupon. The measured porosity values from the XCT samples were between 0.36% and 0.51%. All parts with labeled pores are shown in Figure 9.14-Figure 9.25 in the Appendix.

Table 6.2. Summary of samples containing pores per coupon and coupon XCT porosity

Print	Part	Number of Pores	Porosity
1	1	3047	0.48%
1	2	2957	0.42%
1	3	3015	0.39%
1	4	3305	0.39%
1	5	3157	0.51%
1	6	3029	0.45%
2	7	3638	0.48%
2	8	3128	0.41%
2	9	3846	0.44%
2	10	3805	0.45%
2	11	3000	0.39%
2	12	3187	0.36%

Unlike many previous studies that intentionally induce high porosity to simplify defect detection, this study evaluates localized pore prediction under realistic LPBF processing conditions, with porosity levels ranging from only 0.36% to 0.51%. Under these conditions, pore-related sensor signatures are considerably weaker and more difficult to distinguish from normal process variations, making localized pore prediction substantially more challenging while better representing practical manufacturing environments. As a result, the reported prediction performance more closely reflects the capability of the proposed SFML framework for real-world LPBF applications.

6.3 Individual-Sensor Comparison

Table 6.3 summarizes the performance of the individual sensing modalities for localized pore prediction using process-informed learning. Overall, the validation and test metrics exhibit consistent trends, indicating that the models generalized well to the held-out test coupon. The performance of each sensing modality is evaluated using validation loss, accuracy, AUC, F1 score, precision, and recall.

Table 6.3. Performance of individual sensing modalities with process-informed learning, where py is the pyrometer, coax is the coaxial camera, offax is the off-axis camera, and AE is the acoustic emission.

Sensor*	Val loss	Val accuracy	Val AUC	Val f1	Test accuracy	Test AUC	Test f1	Test Precision	Test Recall
py	0.50	91%	82%	72%	91%	82%	72%	97%	57%
coax	0.37	92%	90%	77%	91%	90%	74%	87%	65%
offax	0.34	92%	93%	79%	91%	93%	77%	79%	76%
AE	0.41	90%	91%	73%	88%	92%	72%	74%	70%

*py: pyrometer, coax: coaxial camera, offax: off-axis camera, AE: acoustic emission

Among the individual sensing modalities, the off-axis branch achieved the best overall performance, with a validation F1 score of 79% and a test F1 score of 77%. It also produced the highest test AUC (93%) and the highest recall (76%), indicating the strongest ability to distinguish pore-containing from non-pore regions while detecting the largest proportion of true pores. The coaxial branch achieved the second-best performance, with a validation F1 score of 77% and a test F1 score of 74%, followed by the AE branch (73% validation F1, 72% test F1) and the pyrometer branch (72% validation F1, 72% test F1).

The confusion matrix for each case is shown in Table 6.4-Table 6.7. Although the pyrometer achieved the highest precision (97%), its recall was the lowest (57%), indicating that it generated very few false-positive predictions but failed to detect a substantial fraction of pore-containing regions. In contrast, the off-axis sensor achieved a more balanced performance, with 79% precision and 76% recall, resulting in the highest F1 score among the individual sensors. The AE branch provided a balanced precision (74%) and recall (70%), whereas the coaxial branch achieved higher precision (87%) at the expense of a lower recall (65%).

The ROC curves for each case are shown in Figure 6.1-Figure 6.4. The pyrometer case produced an ROC AUC of 82%, the coaxial camera produced an ROC AUC of 90%, the off-axis camera produced an ROC AUC of 93%, and the AE case produced an ROC AUC of 92%.

Table 6.4. Confusion Matrix for pyrometer sensors

		Predicted	
		P	N
Actual	P	1884	1421
	N	58	12294

Table 6.5. Confusion Matrix for coaxial camera

		Predicted	
		P	N
Actual	P	2148	1157
	N	321	12031

Table 6.6. Confusion Matrix for off axis camera

		Predicted	
		P	N
Actual	P	2512	793
	N	668	11684

Table 6.7. Confusion Matrix for AE sensor

		Predicted	
		P	N
Actual	P	2313	992
	N	813	11539

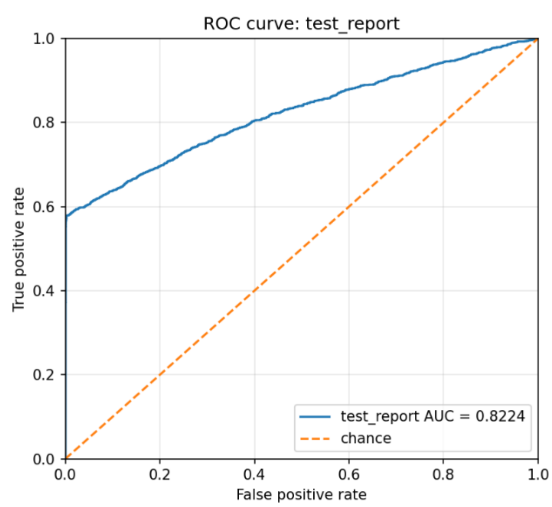

Figure 6.1. ROC curve for pyrometer sensors

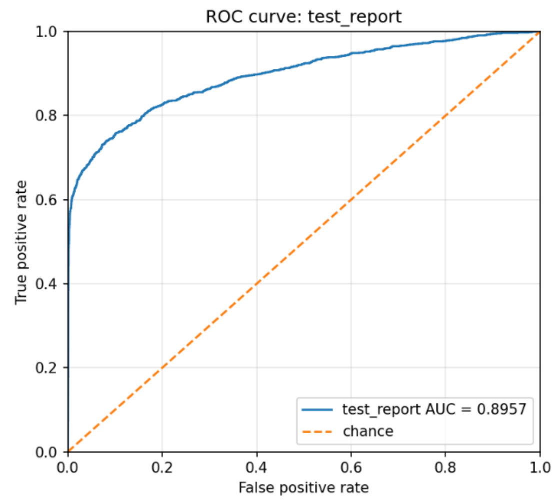

Figure 6.2. ROC curve for coaxial camera

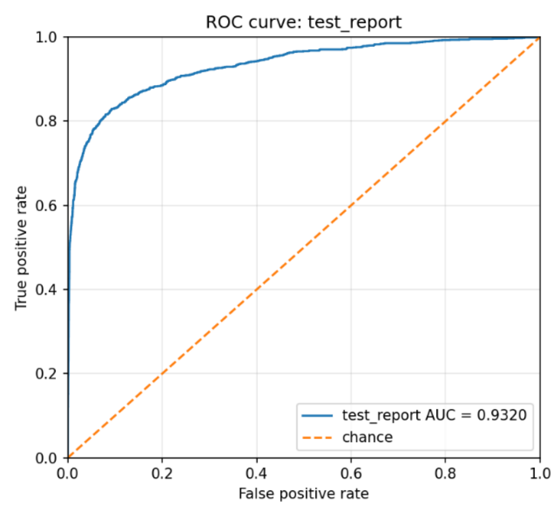

Figure 6.3. ROC curve for off axis camera

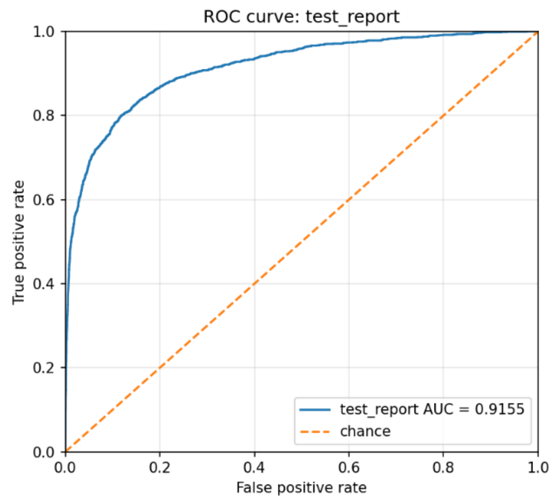

Figure 6.4. ROC curve for AE sensor

6.4 Multi-Sensor Comparison

Table 6.8 summarizes the performance of the multimodal sensor combinations for localized pore prediction. Overall, the validation and test metrics exhibit consistent trends, indicating that the multimodal models generalized well to the held-out test coupon.

Table 6.8. Performance for combined sensors

Sensor	Val loss	Val accuracy	Val AUC	Val f1	Test accuracy	Test AUC	Test f1	Test Precision	Test Recall
coax+py	0.37	92%	90%	77%	91%	88%	76%	88%	67%
offax+py	0.31	94%	94%	83%	93%	93%	82%	90%	75%
ae+py	0.27	94%	95%	84%	93%	95%	82%	88%	77%
offax+coax	0.34	94%	95%	83%	93%	96%	81%	91%	74%
ae+coax	0.27	94%	96%	84%	92%	96%	81%	83%	80%
ae+offax	0.30	94%	96%	84%	92%	96%	81%	85%	77%
ae+py+coax	0.25	95%	96%	86%	93%	96%	84%	86%	81%
ae+py+offax	0.23	95%	97%	87%	93%	96%	84%	86%	81%
all	0.22	95%	97%	88%	94%	96%	85%	85%	85%

Compared with the individual sensing modalities, the two-sensor combinations consistently improved prediction performance. The highest-performing two-sensor configurations were off-axis + pyrometer and AE + pyrometer, both achieving a test F1 score of

82%, representing a 5% improvement over the best individual sensor (77%). These combinations also improved the test AUC to 93–95%, demonstrating that integrating complementary sensing modalities enhanced the model's ability to distinguish pore-containing from non-pore regions.

Further improvements were observed with the three-sensor combinations. Both AE + pyrometer + coaxial and AE + pyrometer + off-axis achieved a test F1 score of 84% with a test AUC of 96%, indicating that incorporating an additional sensing modality further improved localized pore prediction.

The full sensor suite achieved the best overall performance, with a validation accuracy of 95%, validation AUC of 97%, validation F1 score of 88%, test accuracy of 94%, test AUC of 96%, and test F1 score of 85%. The full sensor suite also achieved balanced precision and recall (85% each), demonstrating that integrating all sensing modalities simultaneously reduced both false-positive and false-negative predictions. Overall, the progressive improvement from individual sensors to two-sensor, three-sensor, and four-sensor configurations indicates that pore-related information is distributed across multiple sensing modalities and that these modalities provide complementary information for localized pore prediction.

The corresponding confusion matrices are presented in Table 6.9-Table 6.17, and the ROC curves are shown in Figure 6.5-Figure 6.13.

Table 6.9. Confusion Matrix for coaxial camera and pyrometer test case

		Predicted	
		P	N
Actual	P	2214	1091
	N	302	12050

Table 6.10. Confusion Matrix for off-axis camera and pyrometer test case

		Predicted	
		P	N
Actual	P	2479	826
	N	275	12077

Table 6.11. Confusion Matrix for AE and pyrometer test case

		Predicted	
		P	N
Actual	P	2545	760
	N	347	12005

Table 6.12. Confusion Matrix for off axis camera and coaxial camera test case

		Predicted	
		P	N
Actual	P	2446	859
	N	254	12098

Table 6.13. Confusion Matrix for AE and coaxial camera test case

		Predicted	
		P	N
Actual	P	2644	661
	N	542	11810

Table 6.14. Confusion Matrix for AE and off axis camera test case

		Predicted	
		P	N
Actual	P	2545	760
	N	449	11903

Table 6.15. Confusion Matrix for AE, pyrometer, and coaxial camera test case

		Predicted	
		P	N
Actual	P	2677	628
	N	425	11927

Table 6.16. Confusion Matrix for AE, pyrometer, and off axis camera test case

		Predicted	
		P	N
Actual	P	2690	615
	N	445	11907

Table 6.17. Confusion Matrix for all sensors test case

		Predicted	
		P	N
Actual	P	2806	499
	N	491	11861

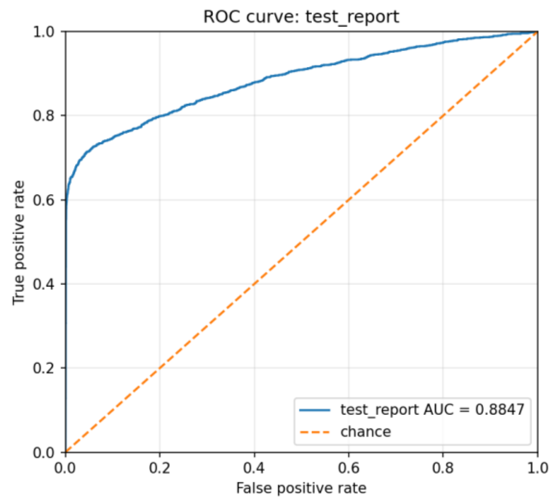

Figure 6.5. ROC Curve for coaxial camera and pyrometer test case

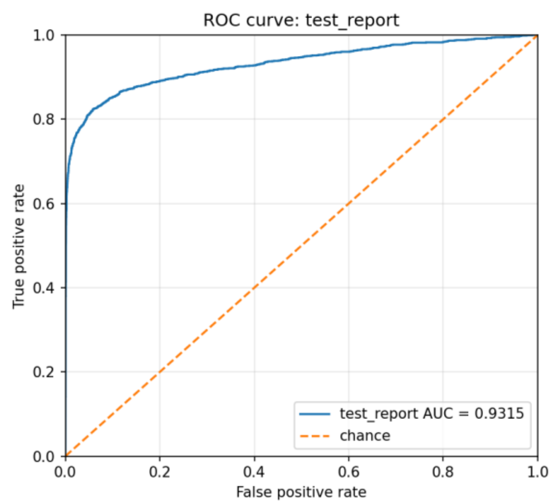

Figure 6.6. ROC Curve for offaxis camera and pyrometer test case

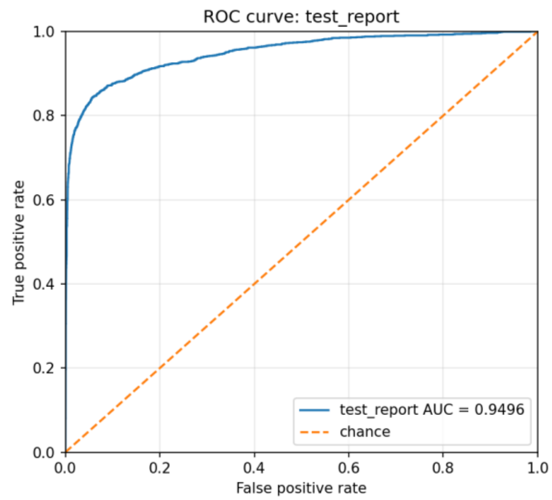

Figure 6.7. ROC Curve for AE and pyrometer test case

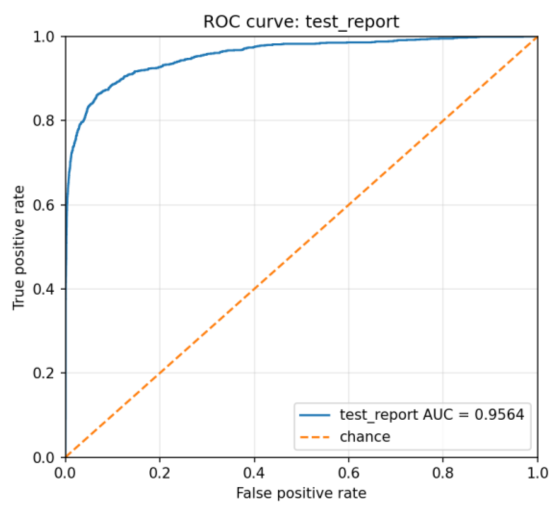

Figure 6.8. ROC Curve for off axis camera and coaxial camera test case

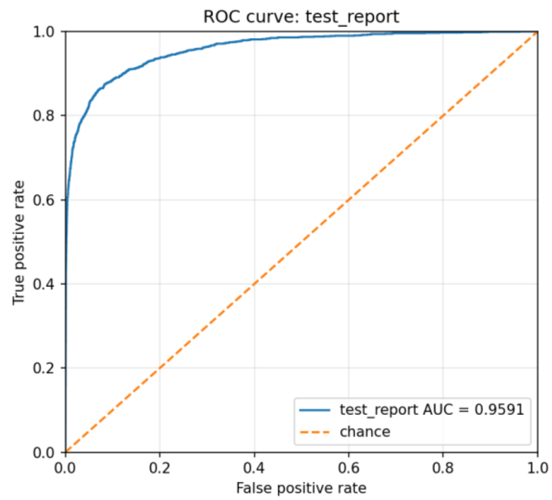

Figure 6.9. ROC Curve for AE and coaxial camera test case

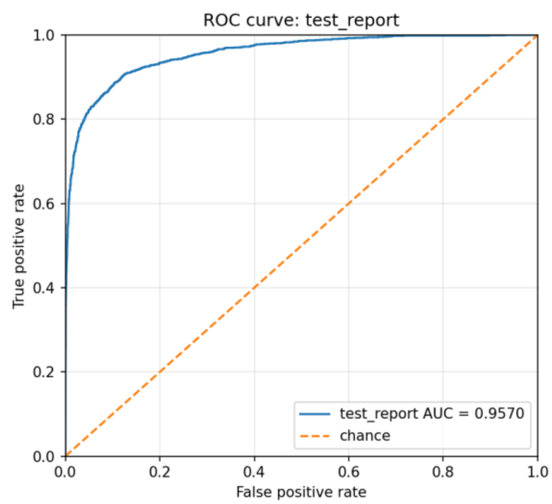

Figure 6.10. ROC Curve for AE and off axis camera test case

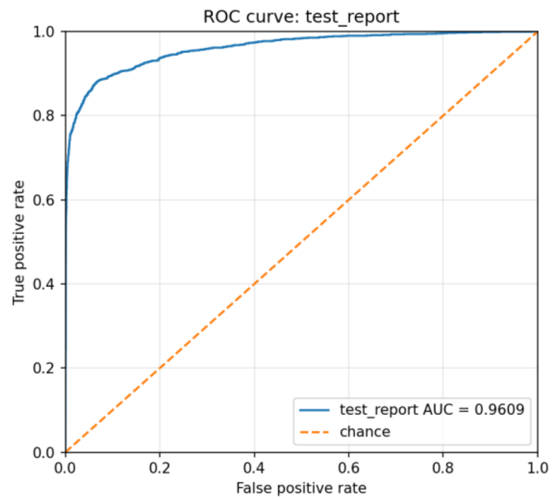

Figure 6.11. ROC Curve for AE, pyrometer, and coaxial camera test case

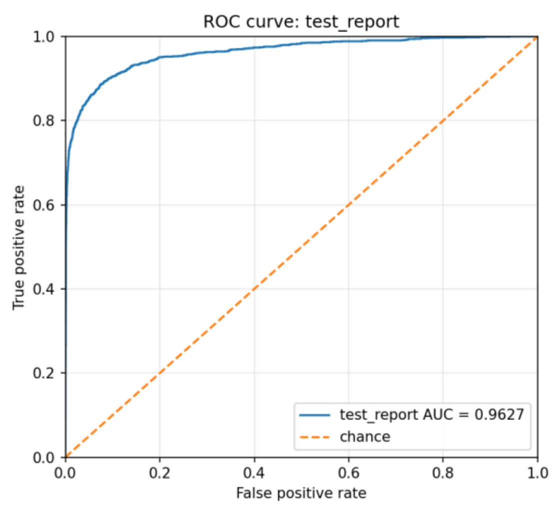

Figure 6.12. ROC Curve for AE, pyrometer, and off axis camera test case

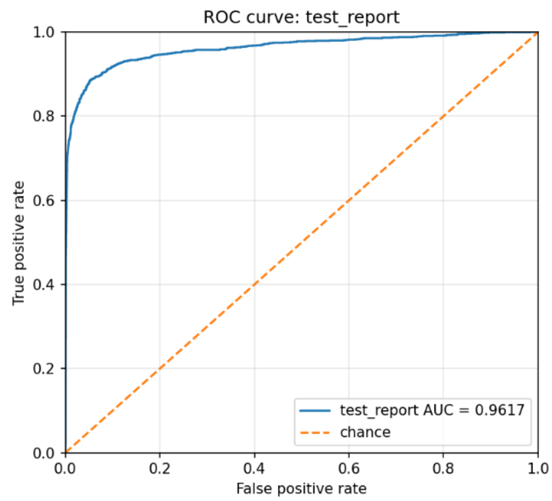

Figure 6.13. ROC Curve for all sensors

6.5 Evaluation of Physics-Informed Learning

Table 6.18 compares the prediction performance of the proposed SFML framework with and without the physics-informed features. All results presented in this section employ the process-informed configuration introduced in Section 3.

Table 6.18. Prediction performance of the full sensor suite with and without physics-informed features

Physcis Informed	Val loss	Val accuracy	Val AUC	Val f1	Test accuracy	Test AUC	Test f1	Test Precision	Test Recall
No	0.22	95%	97%	88%	94%	96%	85%	85%	85%
Yes	0.21	95%	97%	89%	95%	97%	87%	91%	84%

The incorporation of physics-informed features produced the best overall prediction performance. Compared with the baseline model using the full sensor suite, the physics-informed configuration improved the test accuracy from 94% to 95%, the test AUC from 96% to 97%, and the test F1 score from 85% to 87%. Similar improvements were observed for the validation metrics, indicating consistent performance across both the validation and held-out test datasets.

The addition of the physics-informed features increased the test precision from 85% to 91%, while the test recall decreased slightly from 85% to 84%. These results indicate that the physics-informed features substantially reduced false-positive predictions while maintaining a comparable ability to detect pore-containing regions. Overall, the improvement in AUC and F1 score demonstrates that the physics-informed features provide complementary information beyond the multimodal sensor measurements, resulting in more reliable localized pore prediction.

The corresponding confusion matrix is presented in Table 6.19, and the ROC curve is shown in Figure 6.14

Table 6.19. Confusion Matrix for all sensors with and physics-informed features

		Predicted	
		P	N
Actual	P	2776	529
	N	275	12077

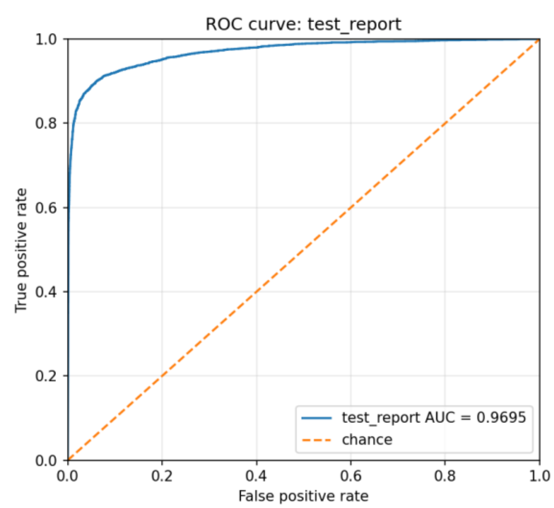

Figure 6.14. ROC Curve for all sensors and physics-informed parameters

6.6 Evaluation of Process-Informed Learning

Section 3.6 introduced the process-informed SoftMax weighting mechanism, which incorporates LPBF process parameters into the sensor-fusion framework to adaptively weight the contributions of different sensing modalities. This ablation isolates the contribution of the process-informed SoftMax weighting used in the preceding sensor and physics-informed comparisons.

Table 6.20 summarizes the prediction performance of the full sensor suite with physics-informed features using the process-informed and non-process-informed configurations. Incorporating process-informed weighting improved the test accuracy from 94% to 95%, the test AUC from 94% to 97%, and the test F1 score from 85% to 87%. Similar improvements were observed for the validation metrics, indicating that the performance gain was consistent across both the validation and held-out test datasets.

Table 6.20. Prediction performance of the full sensor suite and physics-informed features with and without process-informed weighting

Process-Informed	Val Loss	Val Accuracy	Val AUC	Val F1	Test Accuracy	Test AUC	Test F1	Test Precision	Test Recall
No	0.25	95%	96%	85%	94%	94%	85%	92%	78%
Yes	0.21	95%	97%	89%	95%	97%	87%	91%	84%

The process-informed configuration produced a test precision of 91% and a test recall of 84%, whereas the non-process-informed configuration achieved a slightly higher precision (92%) but a substantially lower recall (78%). This result indicates that incorporating process parameters enables the model to identify more pore-containing regions while maintaining a comparable false-positive rate. Consequently, process-informed learning provides complementary process context beyond the multimodal sensor measurements, resulting in improved overall prediction performance.

The corresponding confusion matrix for the non-process-informed configuration is presented in Table 6.21, and the corresponding ROC curve is shown in Figure 6.15.

Table 6.21. Confusion Matrix for all sensors without process-informed weighting

		Predicted	
		P	N
Actual	P	2578	727
	N	215	12137

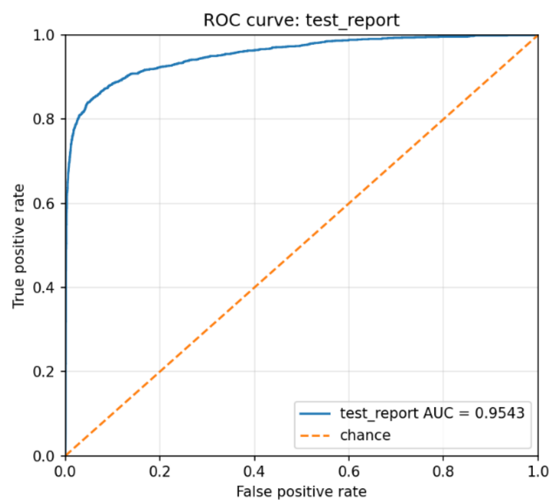

Figure 6.15. ROC curve for all sensors without process-informed weighting

6.7 Fusion-Architecture Comparison

This section compares the proposed PMMSF architecture with two conventional multimodal fusion strategies, input concatenation and head fusion. To isolate the contribution of the fusion architecture, the three methods were first evaluated using the same full multimodal sensor suite without incorporating the proposed physics-informed features or process-informed weighting. The complete proposed framework, PMMSF with physics-informed and process--

informed, was also compared to quantify the additional performance gains provided by these components.

Table 6.22 summarizes the prediction performance of the evaluated fusion architectures. Both PMMSF and head fusion substantially outperformed input concatenation, demonstrating the importance of preserving modality-specific feature extraction for heterogeneous multimodal sensor data. Compared with head fusion, PMMSF consistently achieved higher validation and test performance across the major evaluation metrics, including accuracy, AUC, and F1 score. Specifically, PMMSF improved the validation F1 score from 82% to 87% and the test F1 score from 80% to 83%, while increasing the test AUC from 93% to 95%.

Table 6.22. Comparison of PMMSF with head fusion and input concatenation fusion methods.

Fusion Method	Val Loss	Val Accuracy	Val AUC	Val F1	Test Accuracy	Test AUC	Test F1	Test Precision	Test Recall
Input Concat.	0.34	93%	92%	79%	91%	91%	77%	86%	69%
Head Fusion	0.30	94%	94%	82%	93%	93%	80%	94%	69%
PMMSF	0.27	95%	96%	87%	93%	95%	83%	87%	80%
PMMSF with PIW + PIF*	0.21	95%	97%	89%	95%	97%	87%	91%	84%

*PIW: Process-informed weighting, PIF: Physics-informed features

Although head fusion achieved a higher test precision (94% versus 87%), PMMSF produced a substantially higher recall (80% versus 69%), indicating that it identified a larger proportion of pore-containing regions while missing fewer true pores. These results suggest that PMMSF provides a better balance between pore detection and missed detections, whereas head fusion adopts a more conservative prediction strategy that reduces false-positive predictions at the expense of lower recall.

When the proposed physics-informed features and process-informed weighting were incorporated into PMMSF, the prediction performance further improved to 95% test accuracy, 97% test AUC, and 87% test F1 score, with a precision of 91% and a recall of 84%. These results demonstrate that PMMSF provides a flexible and extensible fusion architecture capable of effectively incorporating additional process and physics information. By integrating these complementary sources of information, the complete SFML framework achieved the highest overall prediction performance.

The corresponding confusion matrices are shown in Table 6.23-Table 6.26 and the corresponding ROC curves are shown in Figure 6.16-Figure 6.19.

Table 6.23. Confusion Matrix for Input Concatenation

		Predicted	
		P	N
Actual	P	2280	1025
	N	371	11981

Table 6.24. Confusion Matrix for Head Fusion

		Predicted	
		P	N
Actual	P	2280	1025
	N	146	12206

Table 6.25. Confusion Matrix for PMMSF

		Predicted	
		P	N
Actual	P	2654	651
	N	408	11944

Table 6.26. Confusion Matrix for PMMSF with process-informed weighting and physics-informed features

		Predicted	
		P	N
Actual	P	2776	529
	N	275	12077

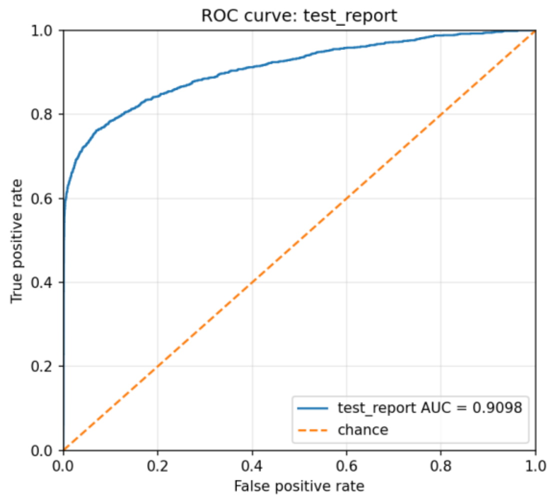

Figure 6.16. ROC curve for Input Concatenation

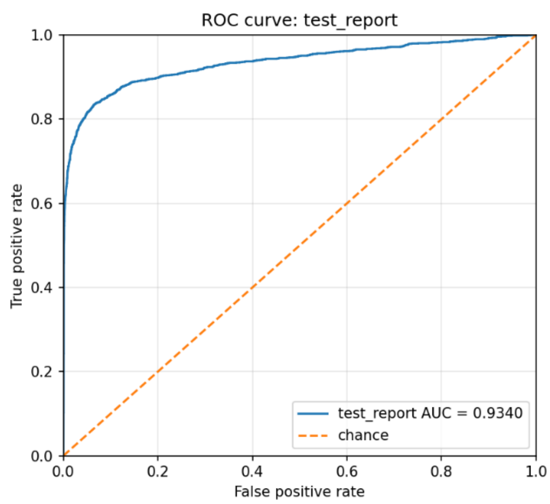

Figure 6.17. ROC Curve for Head Fusion

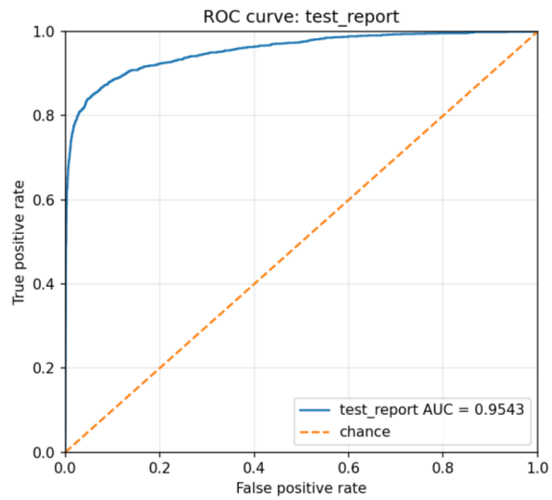

Figure 6.18. ROC Curve for PMMSF

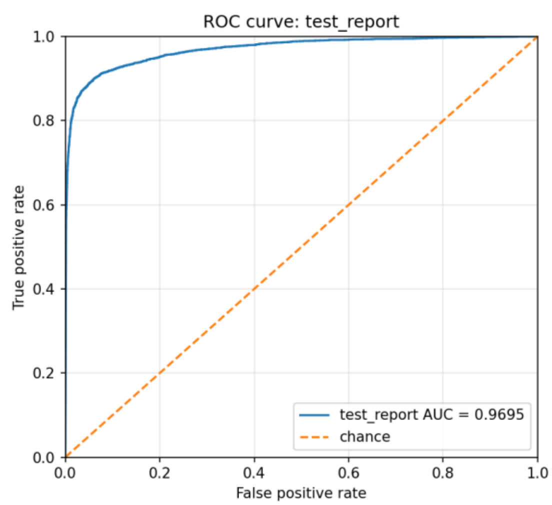

Figure 6.19. ROC curve for PMMSF with process-informed weighting and physics-informed features

6.8 Inference and Alignment Time

This section evaluates the computational efficiency of the proposed SFML framework for localized pore prediction. Table 6.27 summarizes the data alignment, model inference, and total processing time required for each LPBF layer. Table 6.28 summarizes the scanning time and spreading time required for the part print per layer.

Table 6.27. Sensor alignment time with data processing time, model inference time, and total time per layer

Case	Time per Layer
Inference Time	0.05s
Alignment and process Time	2.1s
Total Time	2.2s

Table 6.28. Scan and spreading time for Aconity Mini

Case	Time per Layer
Scan Time	5.1s
Spreading Time	8.5s
Total Time	13.6s

Using the full multimodal sensor suite and a batch size of 512, the average inference time was approximately 0.1 ms per query volume, corresponding to approximately 0.05 seconds to process all query volumes within a single layer. This timing includes only neural network inference and excludes data loading, augmentation, and data type conversion. Sensor registration, physics-informed feature generation, and data preparation required an additional 2.1 s per layer. Consequently, the total processing time was approximately 2.2 seconds per layer. The reported alignment time does not include off-axis camera calibration because this calibration is performed before the printing process.

For the LPBF process used in this study, the powder spreading time and laser scanning time for a single layer were both longer than the total processing time of the proposed framework. Therefore, all data alignment, feature generation, and model inference can be completed before the next layer begins. These results demonstrate that the proposed SFML framework is computationally efficient and capable of supporting near-real-time localized pore prediction during LPBF fabrication.

6.9 Chapter Summary

This chapter presented the quantitative evaluation of the proposed SFML framework for localized XCT-detected pore prediction in LPBF. The dataset consisted of twelve LPBF-fabricated 316L stainless steel coupons containing 39,114 accepted XCT-detected pore components. Individual sensing modalities demonstrated varying prediction capabilities, while progressively integrating complementary sensing modalities consistently improved prediction performance. The full multimodal sensor suite achieved the best performance among the sensor configurations.

The incorporation of physics-informed features and process-informed weighting further improved prediction performance, resulting in the best overall test performance of 95% accuracy, 97% AUC, and 87% F1 score. In addition, the proposed PMMSF architecture outperformed the conventional input concatenation and head fusion strategies, demonstrating the effectiveness of the proposed multimodal fusion framework. Computational efficiency analysis showed that the complete data alignment and inference pipeline can be completed within the processing time of a single LPBF layer, indicating the feasibility of near-real-time localized pore prediction.

The implications of these results, including the contributions of individual sensing modalities, multimodal sensor fusion, physics-informed and process-informed learning, fusion architecture, and the limitations of the proposed framework, are discussed in Chapter 7.

Chapter 7 Discussion

7.1 Overview

This chapter interprets the quantitative results presented in Chapter 6 and discusses their implications for localized porosity prediction in LPBF. The discussion focuses on the contributions of individual sensing modalities, the benefits of multimodal sensing, the effectiveness of the proposed PMMSF architecture, the roles of process-informed and physics-informed learning, the uncertainty associated with XCT-based label generation, and the practical considerations for deploying the proposed SFML framework. Together, these topics provide insight into the factors that govern localized pore prediction and the effectiveness of the proposed framework.

The proposed SFML framework achieved its best performance using the full multimodal sensor suite with process-informed weighting and physics-informed features, producing a test accuracy of 95%, a test AUC of 97%, and a test F1 score of 87%. The corresponding precision and recall were 91% and 84%, respectively, indicating that most predicted pore-containing regions were correct while some true pore-containing regions were still missed. More importantly, the results demonstrate that localized pore prediction benefits from integrating complementary thermal, optical, acoustic, topographic, process-condition, and physics-informed information rather than relying on any single sensing modality. The following sections examine the individual contributions of the sensing modalities, PMMSF, process-informed weighting, and physics-informed features, and discuss their implications for generalized sensor-based prediction and monitoring in complex manufacturing systems.

7.2 Interpretation of Sensor Contributions

The individual-sensor results demonstrate that each sensing modality contains pore-relevant information, although the predictive capability differs among the sensors. Among the individual sensing modalities, the off-axis branch achieved the strongest performance, with a validation F1 score of 79% and a test F1 score of 77%, followed by the coaxial branch (77% validation F1 and 74% test F1). The AE and pyrometer branches achieved lower test F1 scores of 72%, indicating that each sensing modality captures different aspects of the LPBF process.

The superior performance of the off-axis branch is physically reasonable because the off-axis images and photometric-stereo products capture the powder-bed and surface conditions beyond the instantaneous melt-pool region. Variations in powder distribution, surface roughness, remelted geometry, recoater marks, spatter deposition, and previous-layer topography can influence subsequent melting behavior and pore formation. Since the proposed framework incorporates both the current and previous-layer off-axis images, the model can explicitly exploit layer-to-layer surface evolution, providing complementary information that is not directly available from process-time sensing modalities. It is also noted that many pores are located near the part edges and the off-axis camera may provide better geometric context than other sensors.

The coaxial branch also demonstrated strong performance, which is consistent with its direct observation of the laser–material interaction zone. Coaxial images capture variations in melt-pool size, intensity distribution, and morphology that reflect the local processing conditions associated with pore formation. Although these measurements do not directly observe pores, they provide valuable indicators of melt-pool stability and process quality. Incorporating previous-layer coaxial images further enables the model to account for temporal variations in melt-pool behavior.

The AE branch provided moderate performance as an individual sensing modality but contributed more substantially in multimodal configurations. AE measures transient mechanical and acoustic responses generated during the LPBF process, including signals associated with melt-pool instability, spatter, cracking, and other dynamic events. However, these signals are also influenced by wave propagation through the build plate, sensor coupling, and machine vibrations, making them more difficult to interpret independently. Consequently, AE is most effective when combined with complementary optical and topographic information.

The pyrometer branch achieved the lowest individual F1 score despite producing the highest precision. This result suggests that the pyrometer generated relatively few false-positive predictions but failed to identify a larger fraction of pore-containing regions than the other sensing modalities. Because the pyrometer channels were not calibrated as absolute temperature measurements, they should be interpreted as relative thermal-emission signals influenced by emissivity, optical-path response, plume emission, and melt-pool size. Accordingly, their greatest contribution was observed in multimodal sensor fusion, where they complemented the optical, acoustic, and topographic sensing modalities rather than serving as the primary predictor.

Overall, the individual-sensor results demonstrate that no single sensing modality fully characterizes the complex physical mechanisms governing pore formation in LPBF. Instead, each sensor captures complementary information related to the thermal, optical, acoustic, and topographic states of the process, providing the physical basis for the improved prediction performance achieved through multimodal sensor fusion.

7.3 Process-Informed and Physics-Informed Learning

The experimental results demonstrate that both process-informed weighting and physics-informed features provide consistent improvements in localized pore prediction. Although the performance gains are modest compared with those obtained from multimodal sensor fusion, the improvements suggest that incorporating process context and physical knowledge enables the model to interpret sensor measurements more effectively.

The contribution of process-informed weighting can be understood by considering that the same sensor response may have different physical meanings under different process conditions. For example, a pyrometer intensity, melt-pool image feature, AE transient, or surface-topography signature may correspond to different process states depending on laser power and scan speed. Process-informed weighting provides the network with additional process context, allowing it to adapt the relative importance of each sensing modality according to the operating condition. Rather than replacing sensor measurements, the process parameters serve as complementary information that guides multimodal feature fusion.

Similarly, the physics-informed features provide structured physical information that complements the raw sensor measurements. These features should not be interpreted as direct physical state variables because melt-pool depth was not directly measured, the pyrometer channels were not calibrated to absolute temperature, and the photometric-stereo height maps represent relative surface topography rather than absolute geometry. Instead, the physics-informed features serve as physically meaningful proxy variables that help the network interpret multimodal sensor signals in a manner that is more consistent with the underlying LPBF process. However, there is another explanation that should be noted. Physics features include the derivative of the coaxial camera which requires information from the two closest coaxial images. Without physics features only the closest image is fed to the network. This means that additional information is being given to the network that it otherwise would not have access to. The ability to easily compress additional information is perhaps a strength of physics informed features, since including additional images to the network would also increase the training and inference overhead.

The relatively modest performance improvements obtained from both process-informed weighting and physics-informed features also suggest that the synchronized multimodal sensor data remain the primary source of predictive information. Process and physics information alone are insufficient for accurate pore prediction; however, when integrated with multimodal sensing, they provide complementary information that improves the observed predictive performance of the learned representations. These results support the use of process-informed weighting and physics-informed features as complementary sources of information that enhance, rather than replace, multimodal sensor fusion.

7.4 Interpretation of Fusion Performance

The fusion-architecture comparison demonstrates that preserving modality-specific feature learning is essential for multimodal sensor fusion in LPBF. Both PMMSF and head fusion substantially outperformed input concatenation, indicating that heterogeneous sensing modalities should not be fused before meaningful modality-specific features have been extracted. Because the multimodal sensor suite consists of time-series signals, optical images, topographic measurements, and process parameters with different physical meanings and data structures, directly concatenating these heterogeneous inputs limits the network's ability to learn informative feature representations.

Compared with head fusion, PMMSF consistently achieved higher validation and test performance across the major evaluation metrics. Although the performance improvements were modest, they were consistently observed, suggesting that progressively integrating complementary information throughout the network is more effective than performing feature fusion only at the final prediction stage. Furthermore, PMMSF achieved substantially higher recall, indicating an improved ability to identify pore-containing regions while reducing missed detections. This characteristic is particularly important for quality monitoring, where undetected defects are generally more critical than additional false-positive predictions.

The improved performance of PMMSF can be attributed to its progressive fusion strategy. Unlike head fusion, which combines modality-specific features only after each branch has produced a final feature representation, PMMSF allows complementary information to be exchanged repeatedly throughout the feature-learning process while preserving the unique characteristics of each sensing modality. This progressive interaction enables the network to learn richer multimodal representations that better capture the complex relationships among thermal behavior, melt-pool dynamics, acoustic responses, surface conditions, and process parameters associated with pore formation.

An additional advantage of PMMSF is its flexible and extensible architecture. The modular design enables new sources of information to be incorporated without fundamentally changing the overall network structure. In this dissertation, physics-informed features and process-informed weighting were successfully integrated into PMMSF, resulting in the highest overall prediction performance. More broadly, the same architecture can readily accommodate additional sensing modalities or domain-specific knowledge, making PMMSF a flexible foundation for future multimodal monitoring systems.

The multimodal sensor-fusion results further demonstrate that pore-related information is distributed across multiple sensing modalities rather than concentrated in a single sensor. While the best individual sensor achieved reasonable prediction performance, combining complementary thermal, optical, acoustic, topographic, and process information consistently improved prediction accuracy. This observation is consistent with the underlying physics of LPBF, where pore formation results from complex interactions among multiple process phenomena rather than any single measurable variable.

7.5 Registration and Label Uncertainty

Accurate ground-truth label generation is a critical requirement for localized pore prediction because the model learns the relationship between in-situ sensor measurements and post-build XCT observations. Unlike conventional image-classification problems, the labels used in this dissertation are not directly observed during the LPBF process but are generated by registering XCT data to the process-coordinate system. Consequently, the quality of the training labels depends not only on XCT imaging and pore segmentation but also on the accuracy of coordinate registration and spatial alignment.

The proposed center-based labeling strategy provides a practical and reproducible approach for associating localized sensor measurements with individual pores. However, several sources of uncertainty remain. Errors in XCT resolution, pore segmentation, pore-center estimation, coordinate registration, timing synchronization, and layer alignment can shift the estimated pore location relative to the corresponding process location. As a result, pore-containing regions may be incorrectly labeled as negative, while neighboring regions may occasionally be labeled as positive. Such label uncertainty is difficult to eliminate completely and represents an inherent challenge for supervised learning based on post-build inspection.

The query-volume definition also influences the quality of the generated labels. Larger query volumes improve robustness to registration uncertainty but reduce spatial localization, whereas smaller query volumes provide more precise localization at the expense of increased sensitivity to residual alignment errors. The selected query volume in this study represents a compromise between these competing requirements based on the observed pore-size distribution.

Therefore, the reported prediction performance should be interpreted in the context of the uncertainty associated with the generated ground-truth labels. Despite these unavoidable sources of uncertainty, the proposed labeling strategy provides a consistent and practical framework for supervised learning and establishes a foundation for localized pore prediction using multimodal in-situ sensing. Future improvements in XCT resolution, registration accuracy, and label-generation methods are expected to further improve both label quality and prediction performance.

7.6 Practical Deployment Considerations

The proposed SFML framework demonstrates the feasibility of near-real-time localized pore prediction during LPBF. Under the experimental conditions of this study, the complete data-processing pipeline, including data synchronization, preprocessing, and model inference, required approximately 2.2 seconds per layer. This processing time is shorter than the layer processing time of the LPBF process investigated in this study, indicating that layer-wise pore prediction can be completed before the subsequent layer begins.

It should be noted that the reported processing time was obtained using the original implementation without dedicated optimization for deployment. Therefore, the proposed framework represents a proof-of-concept implementation rather than a fully optimized industrial system. Future deployment could benefit from model optimization techniques and hardware-aware acceleration to further reduce computational cost, model size, and inference latency while maintaining prediction performance.

Although the proposed framework demonstrates the computational capability required for layer-wise monitoring, practical deployment will require efficient integration of sensing, data transfer, synchronization, and model inference within the manufacturing workflow. Furthermore, the proposed framework is currently designed for layer-wise monitoring because some sensing information becomes available only after completion of each layer. Achieving continuous real-time monitoring would require advances in sensing hardware, data acquisition, and communication infrastructure.

Another important consideration for industrial deployment is sensor selection. Although the complete multimodal sensor suite achieved the highest prediction performance, reduced sensor configurations may provide a more favorable balance between prediction performance, hardware cost, system complexity, and robustness. Therefore, practical deployment should consider not only prediction accuracy but also sensor cost, calibration requirements, integration complexity, data-management burden, and robustness to sensor failure.

7.7 Limitations

Although the proposed SFML framework demonstrated promising performance for localized pore prediction in LPBF, several limitations should be considered when interpreting the results.

The first limitation is the size and diversity of the dataset. Although the framework was trained using a large number of localized samples, these samples were obtained from only twelve coupons produced in two builds using a single material system (316L stainless steel). Consequently, the generalizability of the proposed framework to different materials, machines, powder batches, scan strategies, and manufacturing environments has not yet been established.

The second limitation is the inherent class imbalance associated with localized pore prediction. Pore-containing regions represent only a small fraction of the build volume, making the prediction task highly imbalanced. Although balanced sampling improves training stability, it does not reflect the natural class distribution encountered during practical deployment. Therefore, the reported performance should be interpreted using multiple evaluation metrics rather than a single metric.

The third limitation is the uncertainty associated with XCT-based label generation. The ground-truth labels depend on XCT resolution, pore segmentation, coordinate registration, and query-volume definition, all of which introduce unavoidable uncertainty into the learning labels. Consequently, the reported prediction performance reflects not only the capability of the proposed SFML framework but also the quality of the generated ground-truth labels.

The fourth limitation concerns the sensing system. The pyrometer channels were not calibrated as absolute temperature measurements and therefore provide relative thermal-emission information rather than true temperature. In addition, the measured AE signals are influenced by wave propagation, sensor coupling, and structural vibrations. Although these sensing limitations did not prevent effective multimodal sensor fusion, they should be considered when interpreting the learned feature representations.

The fifth limitation is the limited process window investigated in this study. The experiments varied laser power and scan speed while maintaining other processing parameters, such as layer thickness, hatch spacing, and spot size, at fixed values. Consequently, the proposed framework has not yet been validated across the full range of LPBF processing conditions.

Overall, the proposed framework demonstrated promising localized pore-prediction capability under the investigated conditions; however, additional validation across broader manufacturing conditions is required before its general applicability can be fully established.

7.8 Chapter Summary

This chapter discussed the experimental results and their implications for localized pore prediction in LPBF. The individual sensor analysis demonstrated that each sensing modality provides complementary information related to the thermal, optical, acoustic, topographic, and process states of the LPBF process, establishing the physical basis for multimodal sensor fusion.

The discussion further showed that process-informed weighting and physics-informed features provide complementary process knowledge that enhances multimodal sensor fusion, while the proposed PMMSF architecture enables more effective integration of heterogeneous sensing modalities through progressive feature fusion. The results also highlighted the importance of accurate XCT-to-process registration and reliable label generation for supervised learning, as uncertainties in the generated ground-truth labels directly influence the achievable prediction performance.

Finally, the practical considerations and limitations of the proposed framework were discussed. The results demonstrate the feasibility of near-real-time layer-wise localized pore prediction while identifying several challenges related to dataset diversity, label uncertainty, and validation across broader manufacturing conditions. Overall, the proposed SFML framework provides an effective and extensible approach for multimodal in-situ monitoring of LPBF processes and establishes a foundation for future intelligent quality monitoring in metal additive manufacturing. The next chapter summarizes the major contributions of this dissertation and presents recommendations for future research.

Chapter 8 Conclusion and Future Study

8.1 Dissertation Conclusions

This dissertation developed and validated a Sensor Fusion Machine Learning (SFML) framework for localized prediction of X-ray Computed Tomography (XCT)-detected porosity in laser powder bed fusion (LPBF). The proposed framework addressed the fundamental challenge of linking heterogeneous in-situ sensor measurements with post-build XCT observations by integrating multimodal sensing, process-coordinate synchronization, localized ground-truth generation, process-informed weighting, physics-informed features, and the proposed Parallel Multi-Layer Multimodal Sensor Fusion (PMMSF) architecture.

The first conclusion is that heterogeneous LPBF sensor data can be successfully synchronized within a common process-coordinate framework for localized machine learning. The proposed data preparation workflow effectively aligned process data, pyrometry, acoustic emission, coaxial imaging, off-axis imaging, and photometric-stereo measurements, enabling multimodal sensor data to be associated with specific process locations throughout the build.

The second conclusion is that XCT-derived pore information can be transformed into localized learning labels through pore segmentation, coordinate registration, and process-centered query-volume labeling. This localized labeling strategy enables supervised learning at the process-location level rather than at the part level, providing a practical framework for integrating post-build inspection with in-situ process monitoring.

The third conclusion is that multimodal sensing significantly improves localized pore prediction compared with individual sensing modalities. The best individual sensing modality achieved a test accuracy of 91% and a test F1 score of 77%, whereas the complete multimodal sensor suite improved the performance to 94% test accuracy and 85% test F1 score. These results demonstrate that thermal, optical, acoustic, topographic, and process information provide complementary observations that collectively enable more reliable defect prediction.

The fourth conclusion is that the proposed PMMSF architecture provides a more effective multimodal fusion strategy than conventional fusion methods. PMMSF consistently outperformed both input concatenation and head fusion, improving the test F1 score from 77% (input concatenation) and 80% (head fusion) to 83% while achieving the highest test AUC of 95%. These results demonstrate that preserving modality-specific feature learning while progressively integrating complementary information is more effective than conventional fusion approaches for heterogeneous LPBF sensing data.

The fifth conclusion is that incorporating process-informed weighting and physics-informed features further enhances localized pore prediction by providing complementary manufacturing knowledge beyond the measured sensor signals. Integrating these components into PMMSF further improved the prediction performance to 95% test accuracy, 97% test AUC, and 87% test F1 score, demonstrating the effectiveness of combining multimodal sensing with process knowledge and physics-informed information.

Overall, this dissertation establishes a comprehensive framework for localized pore prediction in LPBF by integrating multimodal sensing, process-coordinate synchronization, localized XCT-based ground-truth generation, process-informed weighting, physics-informed features, and the proposed PMMSF architecture. Beyond localized pore prediction, the proposed SFML framework provides a general methodology for integrating heterogeneous sensing, manufacturing knowledge, and physics-informed information into intelligent monitoring systems for advanced manufacturing. The concepts and methodologies developed in this dissertation establish a foundation for future in-situ quality monitoring, intelligent process optimization, and autonomous manufacturing systems.

8.2 Contributions

This dissertation makes six primary contributions to localized pore prediction and multimodal sensor fusion for LPBF.

First, the proposed SFML framework was developed for localized prediction of XCT-detected porosity in LPBF. The framework integrates multimodal sensing, process-coordinate synchronization, localized ground-truth generation, process-informed weighting, physics-informed features, and PMMSF into a unified workflow for in-situ quality monitoring.

Second, a multimodal data-acquisition and process-coordinate synchronization workflow was developed to organize heterogeneous sensing modalities within a common spatial and temporal framework. Process data, pyrometry, AE, laser-trigger signals, coaxial images, off-axis images, and photometric-stereo height maps were synchronized and stored in a unified HDF5 dataset, enabling heterogeneous sensor measurements to be associated with the same process location and layer.

Third, a localized XCT-based ground-truth generation workflow was developed for supervised learning. The proposed workflow combines XCT pore segmentation, coordinate registration, and process-centered query-volume labeling to transform post-build XCT observations into localized pore labels, establishing a practical methodology for linking post-build inspection with in-situ process monitoring.

Fourth, a multimodal feature-engineering framework was developed to support localized pore prediction. The framework integrates raw sensor measurements with process-informed weighting and physics-informed features, enabling complementary manufacturing knowledge to be incorporated into multimodal learning while preserving the original sensor information.

Fifth, the proposed PMMSF architecture was developed for heterogeneous multimodal sensing. Unlike conventional fusion strategies, PMMSF preserves modality-specific feature extraction while progressively integrating complementary information across multiple network stages. The modular architecture also provides the flexibility to incorporate additional sensing modalities and domain-specific information, making it readily extensible to future multimodal monitoring systems.

Finally, a comprehensive experimental framework was established to systematically evaluate individual sensing modalities, multimodal sensor combinations, fusion architectures, process-informed weighting, and physics-informed features. This framework provides a benchmark for evaluating multimodal AI-based monitoring systems and demonstrates the effectiveness of integrating heterogeneous sensing, manufacturing knowledge, and progressive multimodal sensor fusion for LPBF.

8.3 Future Study

Several research directions are recommended to further improve the proposed SFML framework and expand its applicability to industrial LPBF systems.

First, future work should validate the proposed framework using larger and more diverse datasets. Although the present study demonstrated the feasibility of localized pore prediction using twelve coupons fabricated from 316L stainless steel, broader validation is required to establish the generalizability of the framework. Future studies should include additional builds, machines, materials, scan strategies, and process windows to evaluate part-to-part, build-to-build, and machine-to-machine robustness. Extending the framework to materials such as Ti-

6Al-4V, nickel-based superalloys, aluminum alloys, and tool steels will also help distinguish material-specific characteristics from generalizable multimodal sensing principles.

Second, future research should further improve the physical representation of the monitoring framework. Additional process conditions, including hatch spacing, layer thickness, scan strategy, gas flow, powder reuse, and build geometry, should be investigated to better understand their influence on multimodal sensor fusion. Furthermore, calibrated pyrometry and direct melt-pool-depth measurements should be incorporated to provide more physically meaningful features and improve both prediction performance and model interpretability.

Third, future work should focus on practical deployment of the proposed framework. Although the current implementation demonstrated the feasibility of near-real-time layer-wise prediction, the model was not specifically optimized for deployment. Model-compression techniques, such as pruning, quantization, and knowledge distillation, together with optimized inference pipelines and edge-computing platforms, could significantly reduce computational cost, model size, and inference latency while maintaining prediction performance. These developments would facilitate practical implementation of the proposed framework in industrial LPBF systems.

Fourth, improvements in ground-truth label generation should be investigated. Future studies should quantify uncertainties associated with XCT resolution, pore segmentation, coordinate registration, timing synchronization, layer alignment, and query-volume definition. Probabilistic labeling and uncertainty-aware learning may further improve model robustness. In addition, higher-resolution XCT imaging and optimized scanning conditions could provide more accurate pore characterization and improve the quality of the ground-truth labels generated.

Finally, the proposed framework should be extended beyond localized pore prediction. The same multimodal sensing and sensor-fusion framework could be applied to other LPBF quality characteristics, including microstructure, grain morphology, surface roughness, lack-of-fusion defects, cracking, residual stress, and local mechanical properties. Ultimately, integrating localized defect prediction with microstructural evolution and mechanical-property prediction could establish a comprehensive in-situ quality-monitoring framework capable of supporting qualification and intelligent process control in metal additive manufacturing.

Chapter 9 Appendix

9.1 Custom Software

This appendix contains screenshots of the custom software developed for this project. The primary custom-software interfaces and processing workflows are shown in Figure 9.1-Figure 9.13.

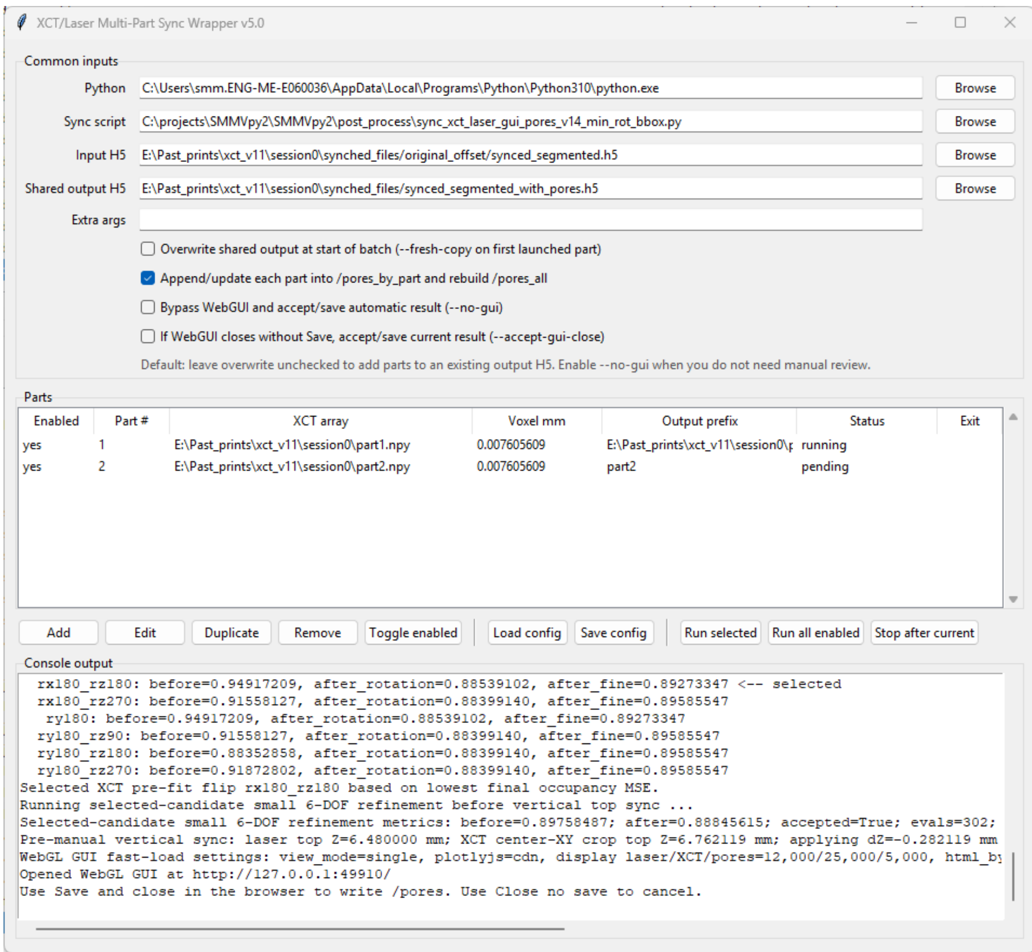

Figure 9.1. Wrapper for XCT part to sensor data synchronization. Multiple parts can be loaded and the registered pores can be saved to an output h5 file.

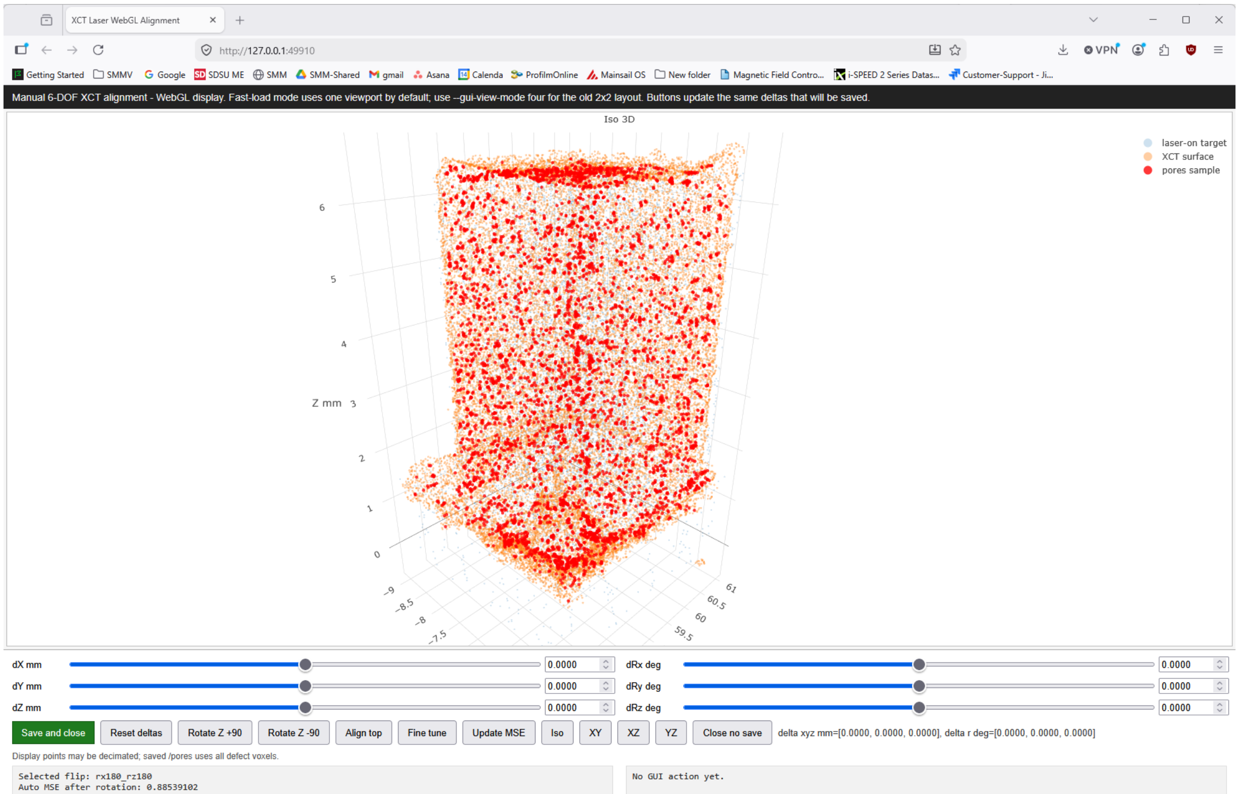

Figure 9.2. XCT volume and pore alignment with laser trigger signals. Manual alignment can be facilitated as well within the GUI and saved.

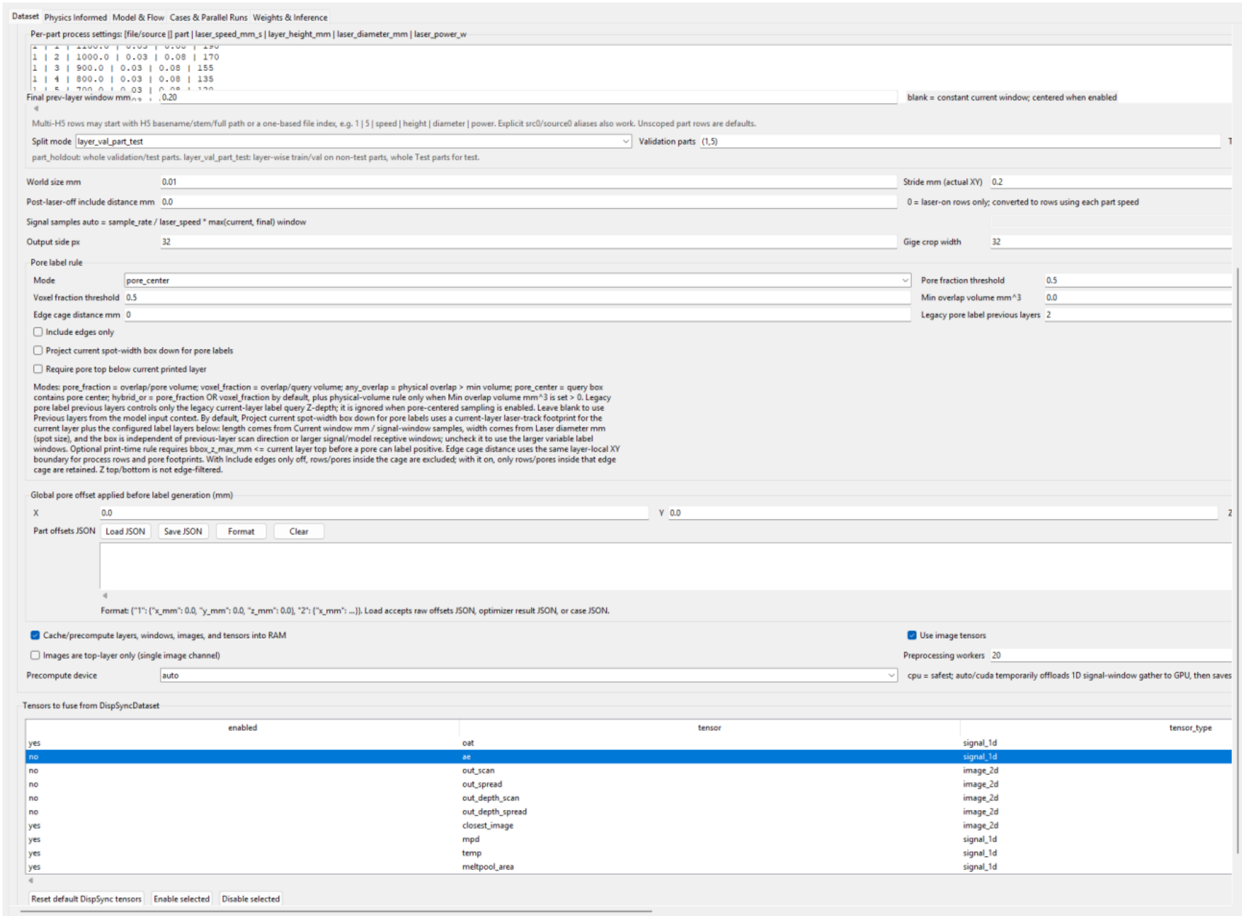

Figure 9.3. Fusion Studio v192 case setup. The query volume, sensor selection, offset, pore matching mode, stride, validation mode can all be specified.

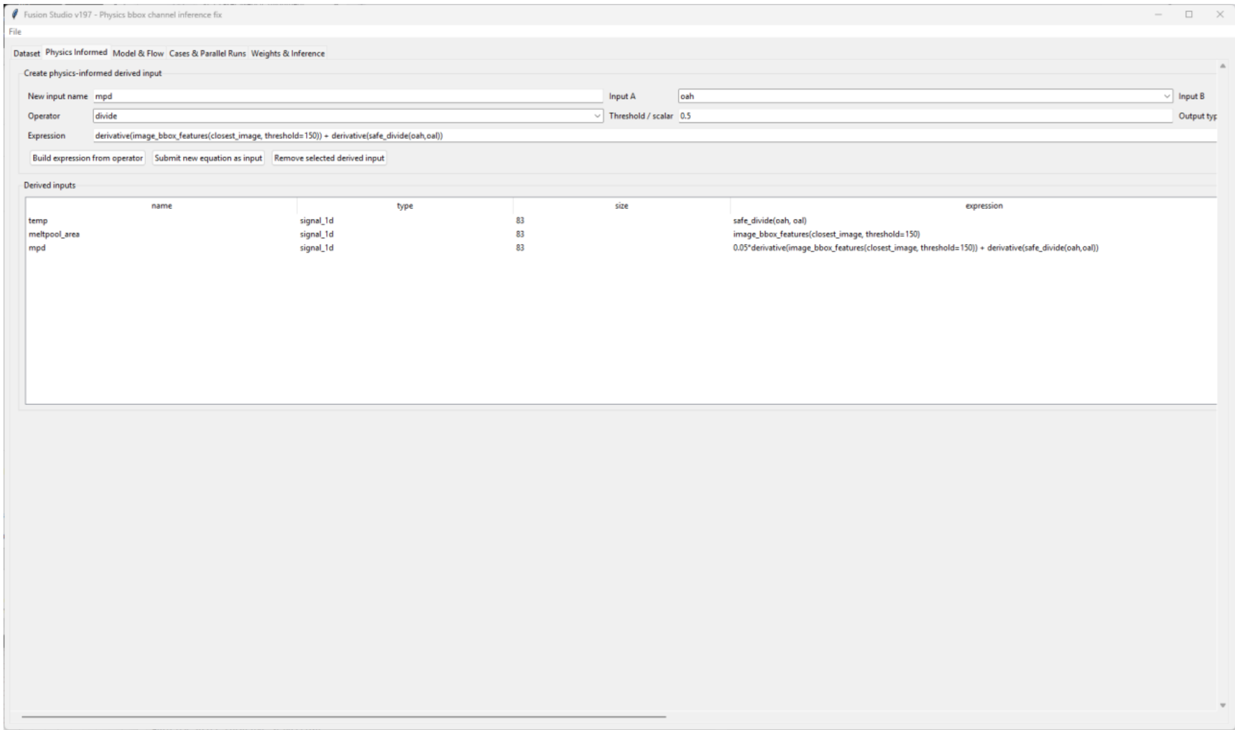

Figure 9.4. Physics and feature extraction window of Fusion Studio v192 part 1. The figure shows how to input custom equations for informing physics.

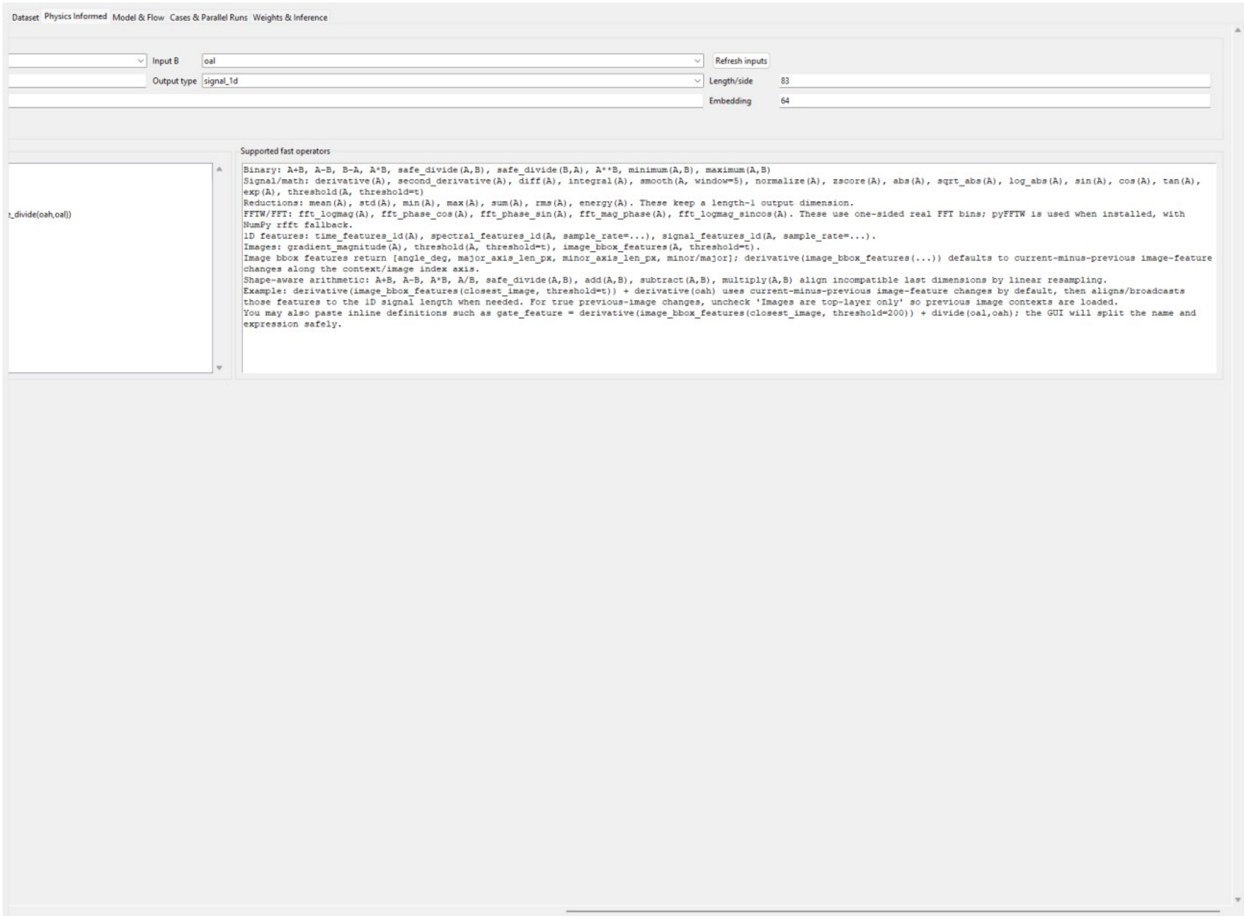

Figure 9.5. Physics and feature extraction window of Fusion Studio v192 part 2. Various signal processing functions as well as standard mathematical operators can be applied.

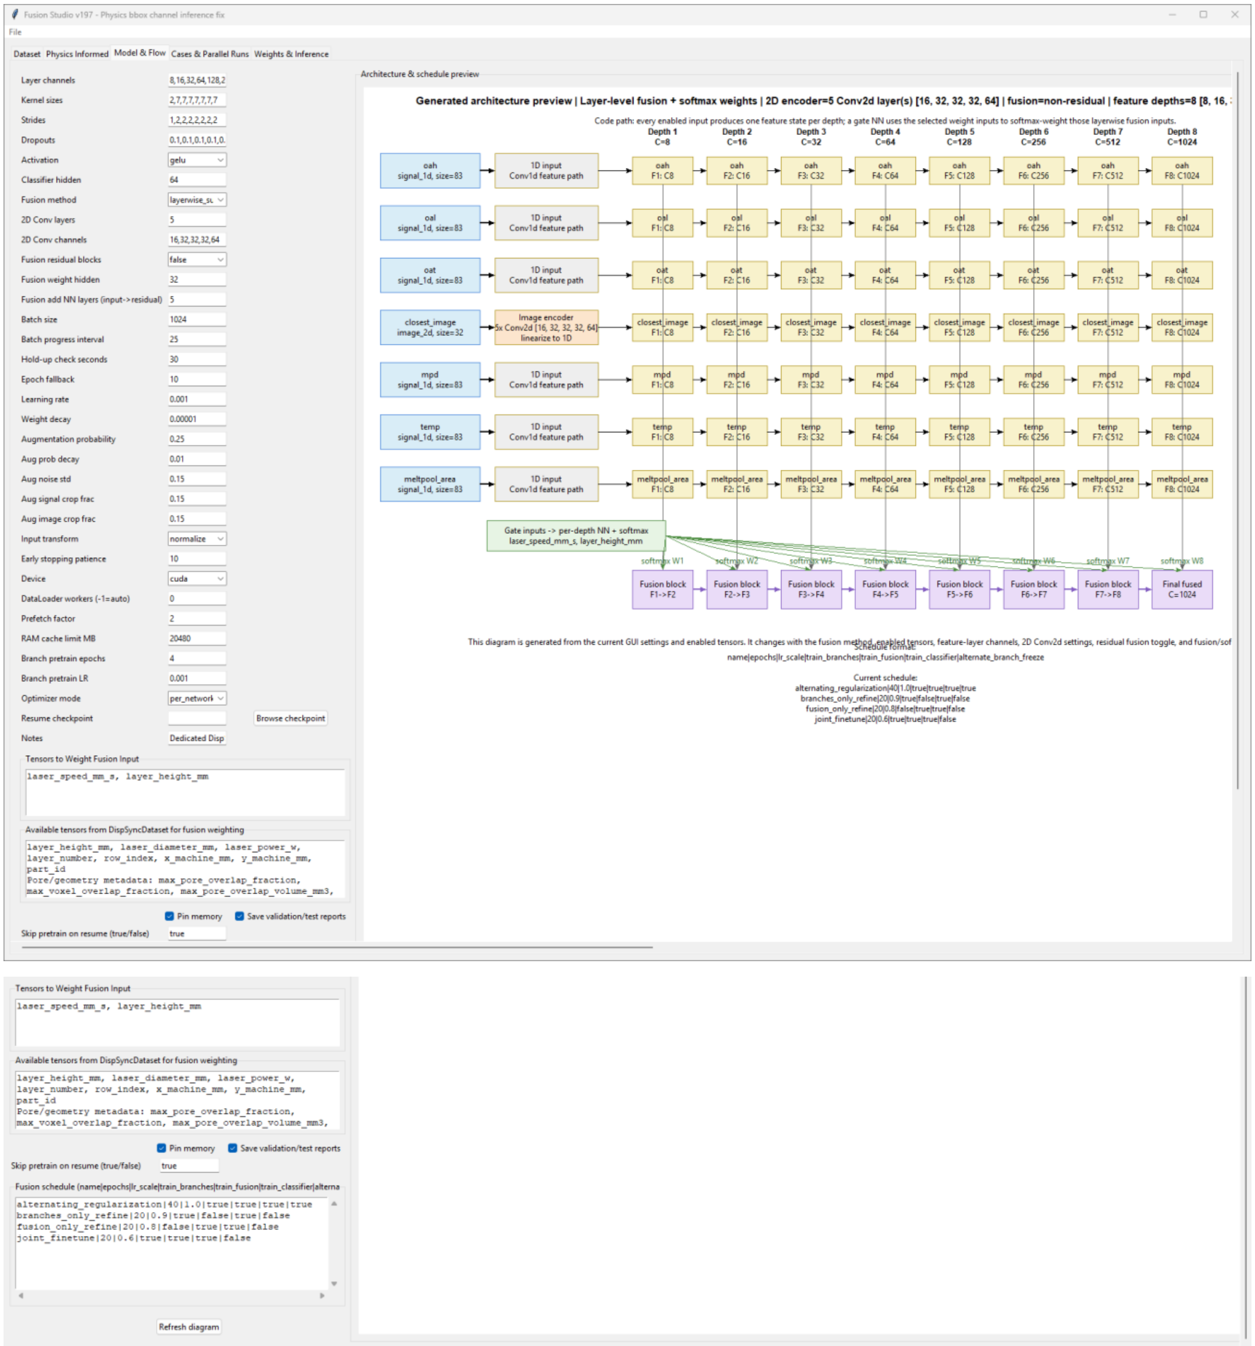

Figure 9.6. Fusion Studio v192 model design and machine learning setup (top). Additional tensor fusion weight can be added by specifying in the physics-informed window and typing them in here (bottom) the fusion schedule can also be modified by updating the window parameters

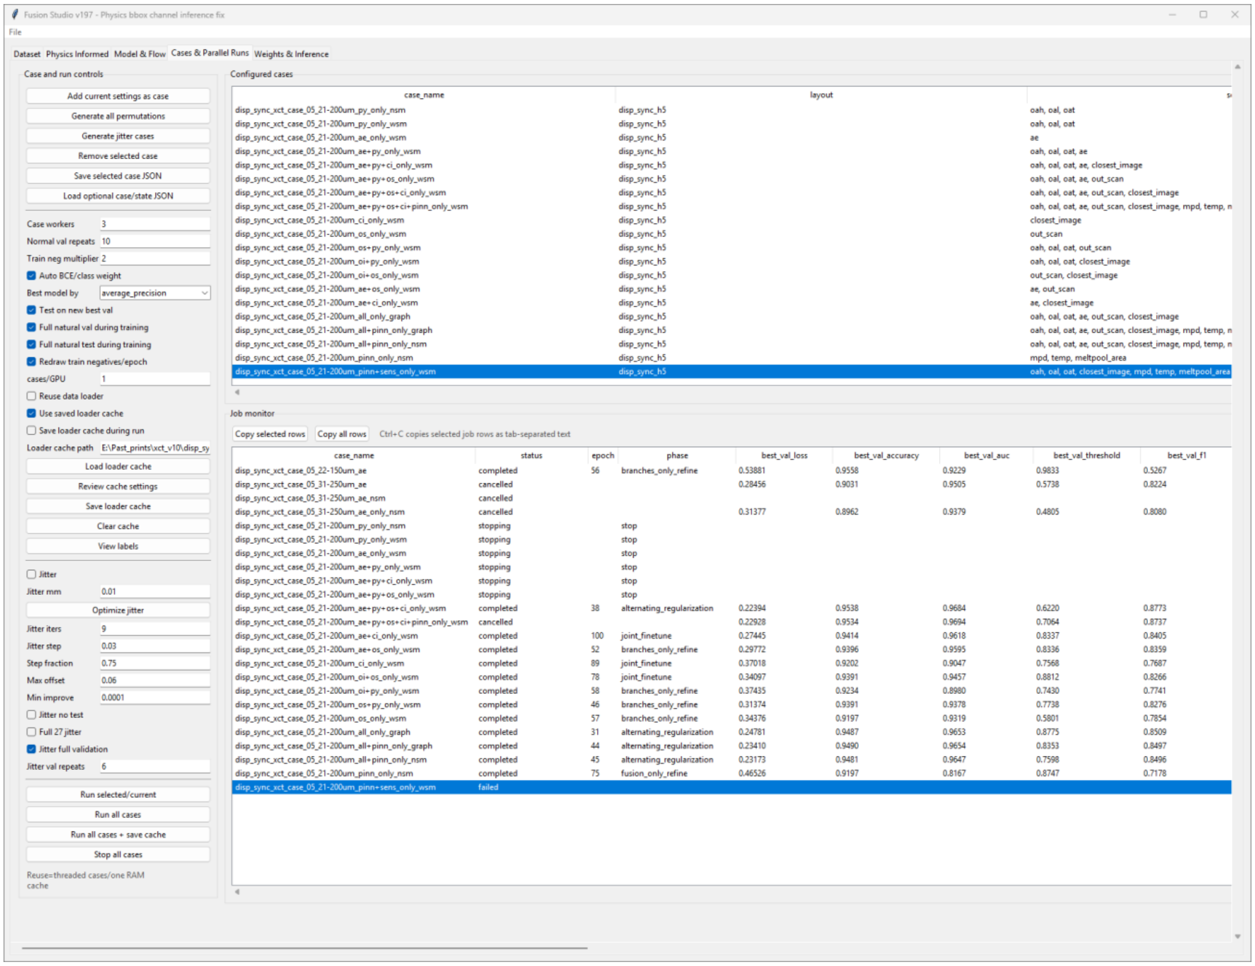

Figure 9.7. Fusion Studio training monitor. Cases are added to the top and then after submitting are added to the queue and shown in the training monitor.

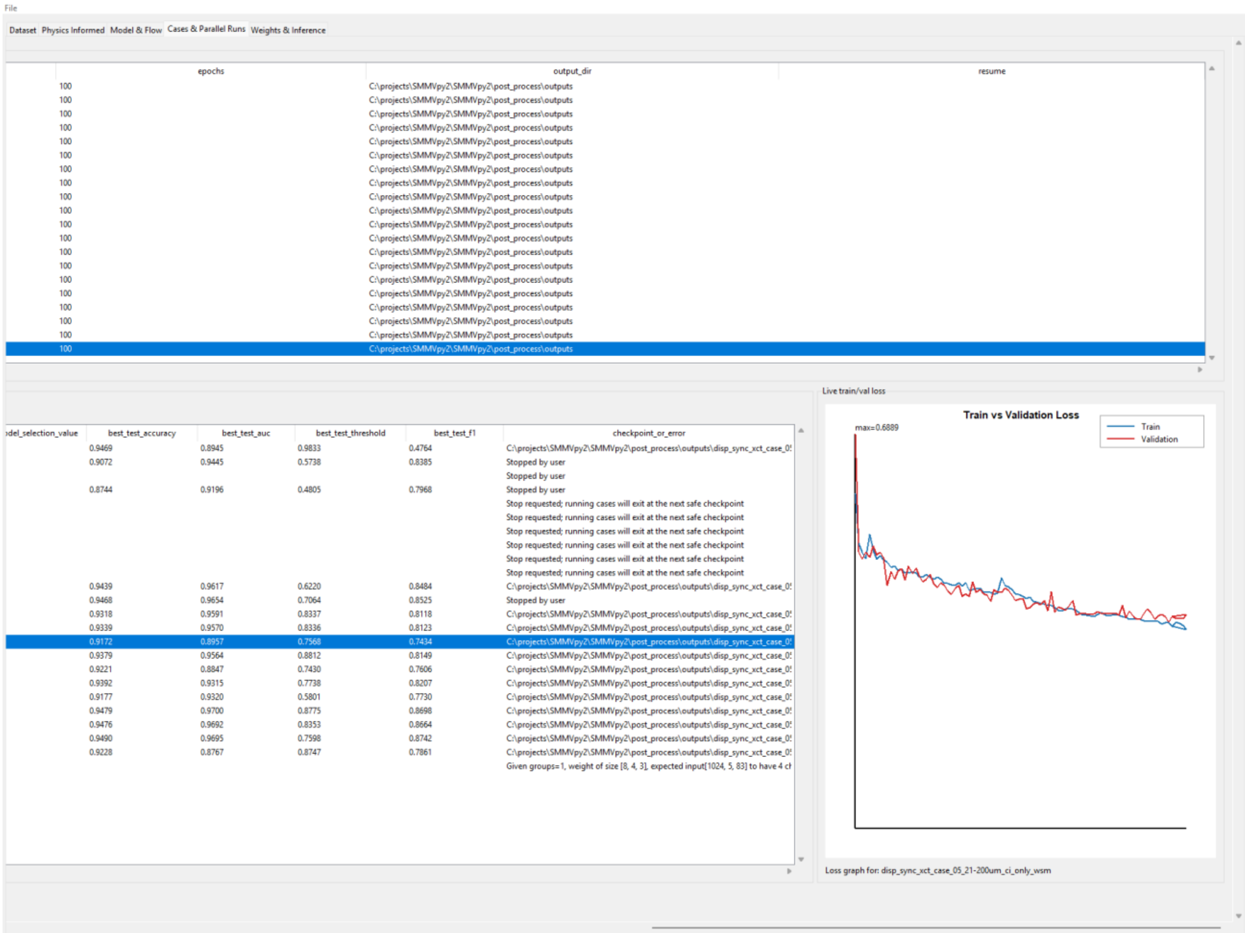

Figure 9.8. Fusion Studio v192 training monitor pt2. Each case displays a real-time loss function for the training as well as the validation and testing results.

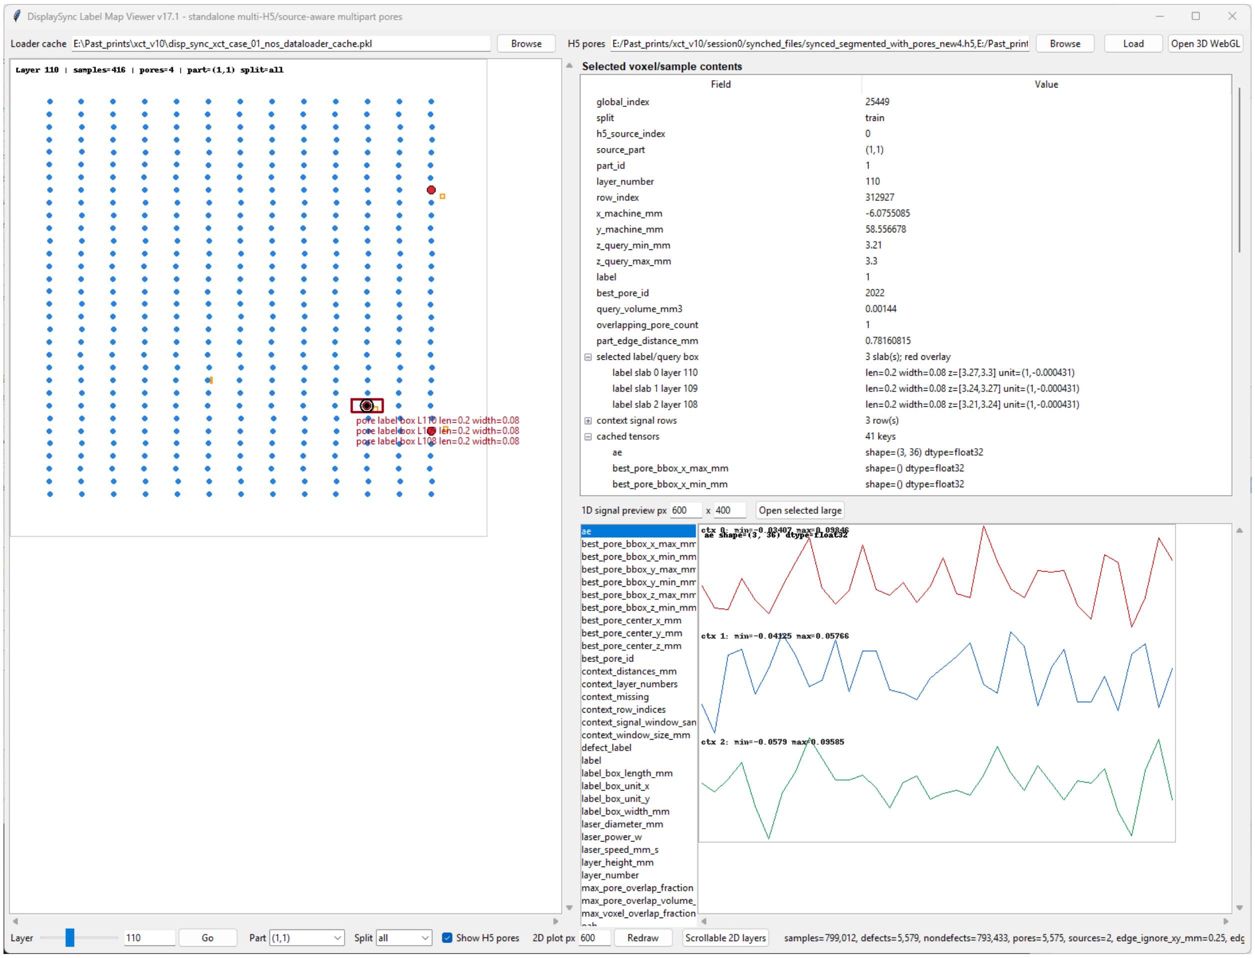

Figure 9.9. Fusion Studio v192 data viewer. This allows the user to view what datapoints are marked as pores and what the data looks like. The pore location is also marked as a green dot.

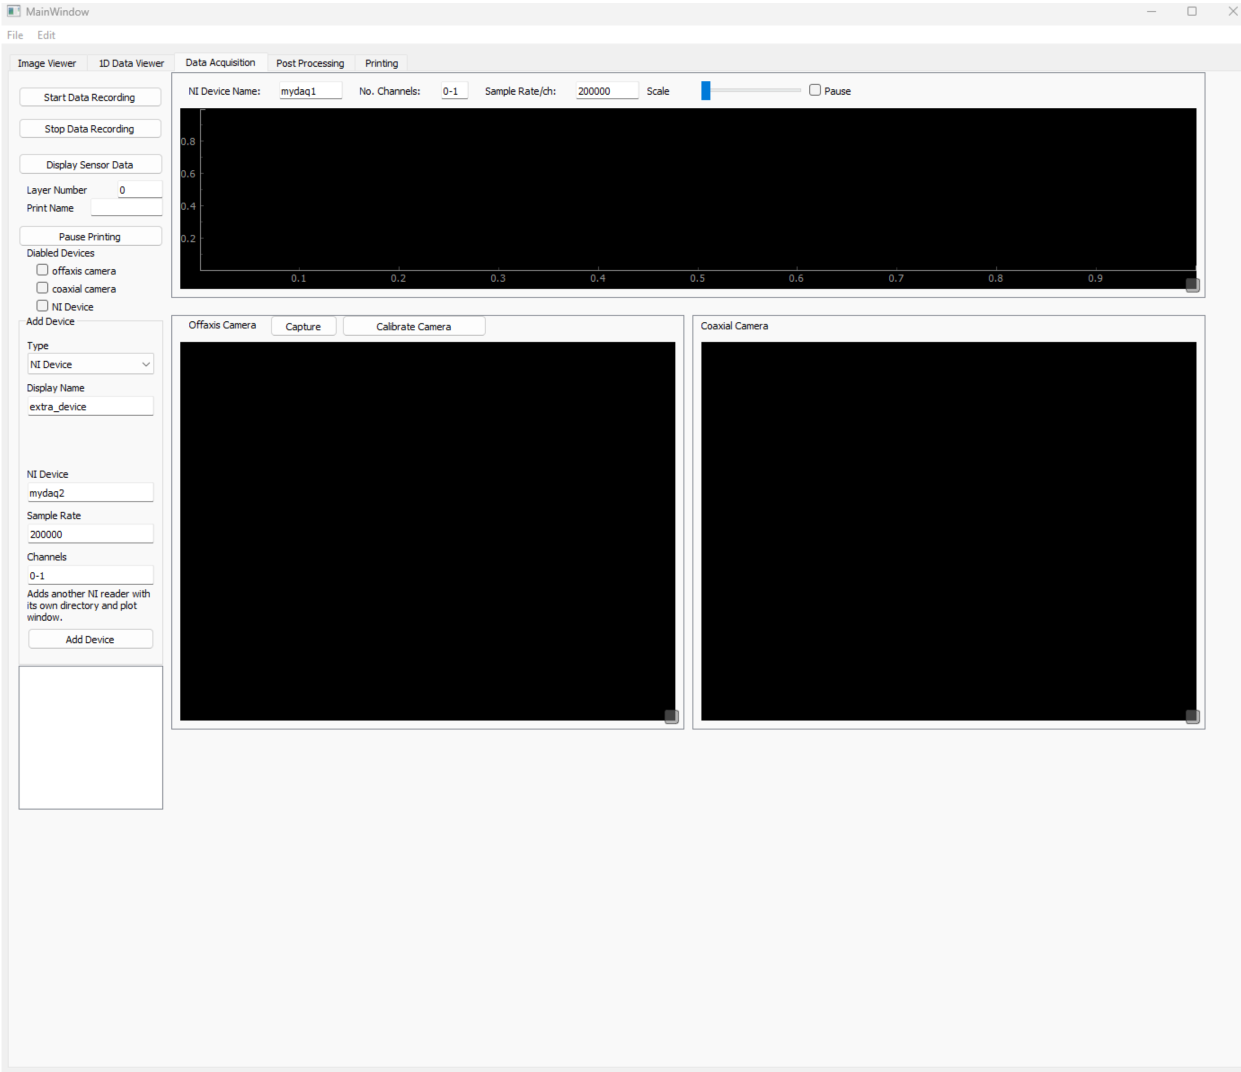

Figure 9.10. SMMVpy2 data collection screen. Real-time data is shown and collected with dedicated processes.

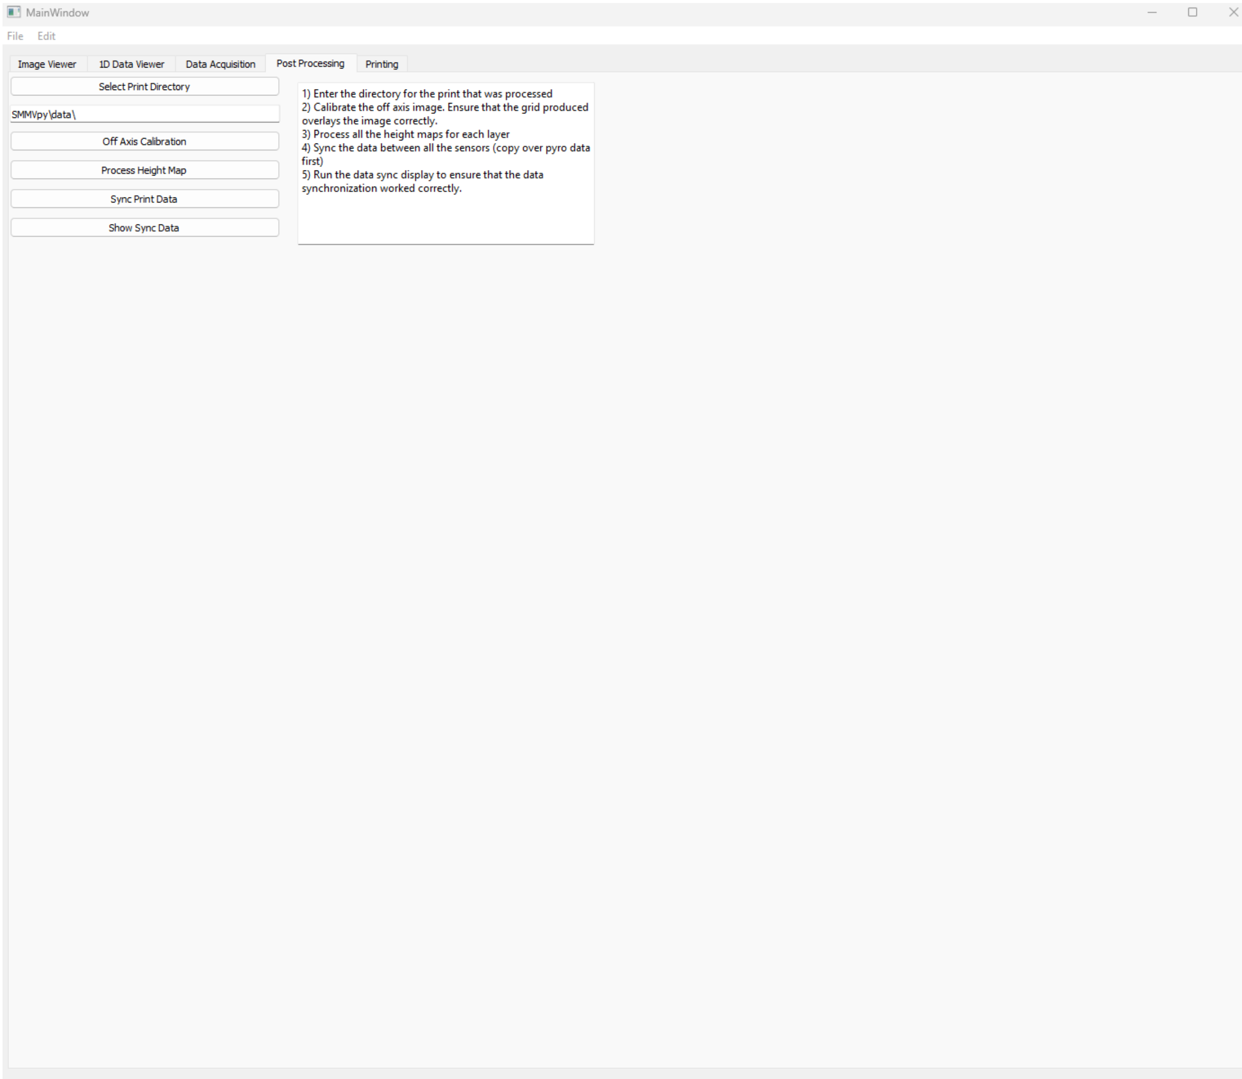

Figure 9.11. SMMVpy2 post processing screen. After a print is completed, the user selects each processing step to perform camera calibration, height map processing and sensor alignment.

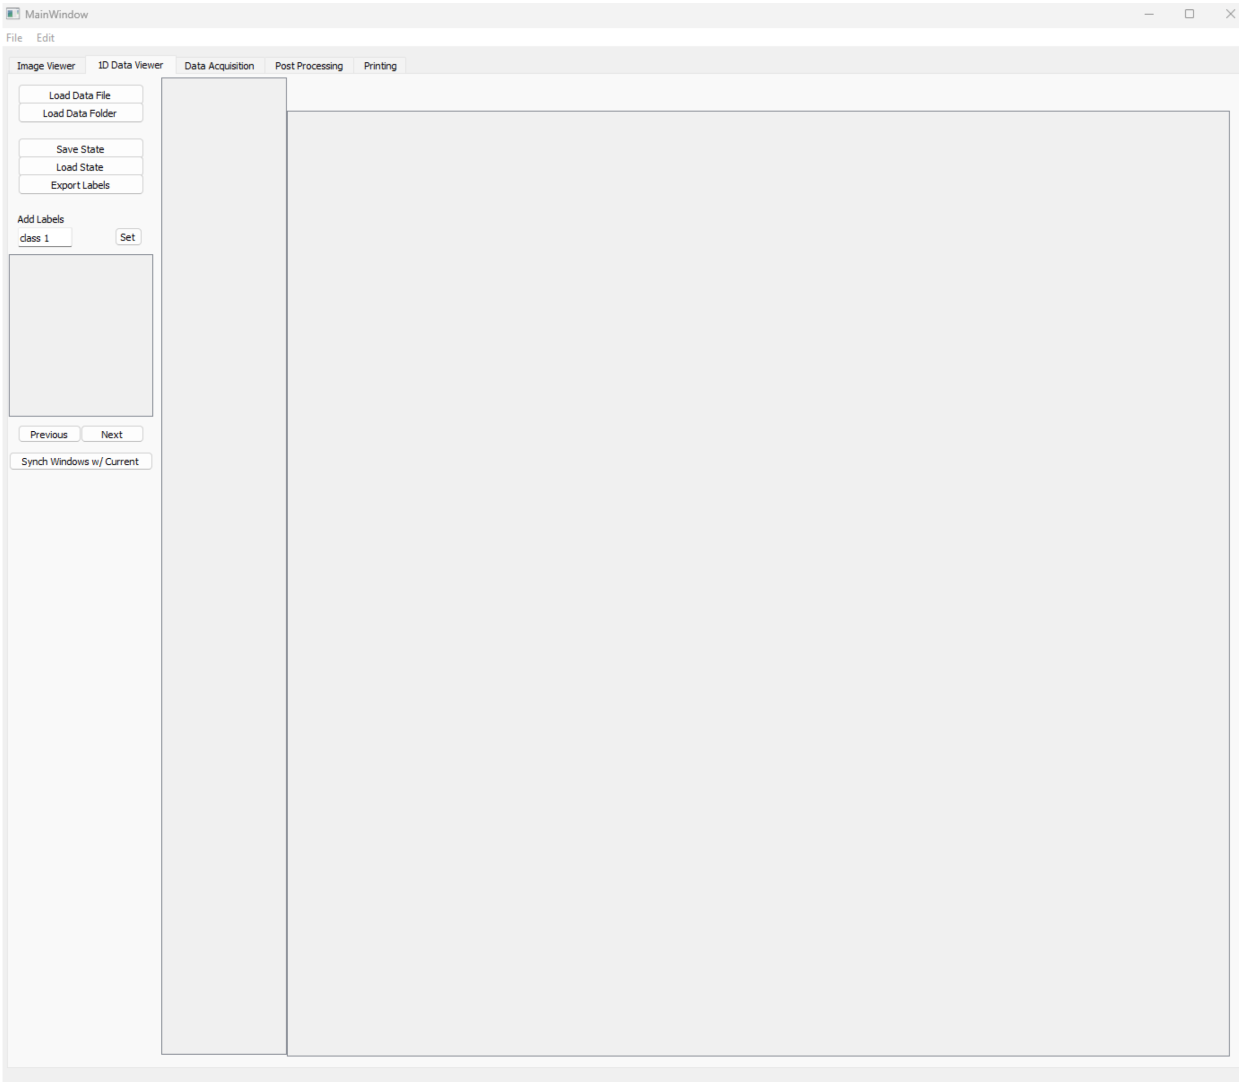

Figure 9.12. SMMVpy2 1D data viewer. 1D data can be viewed, and label classes can be added for manual labeling

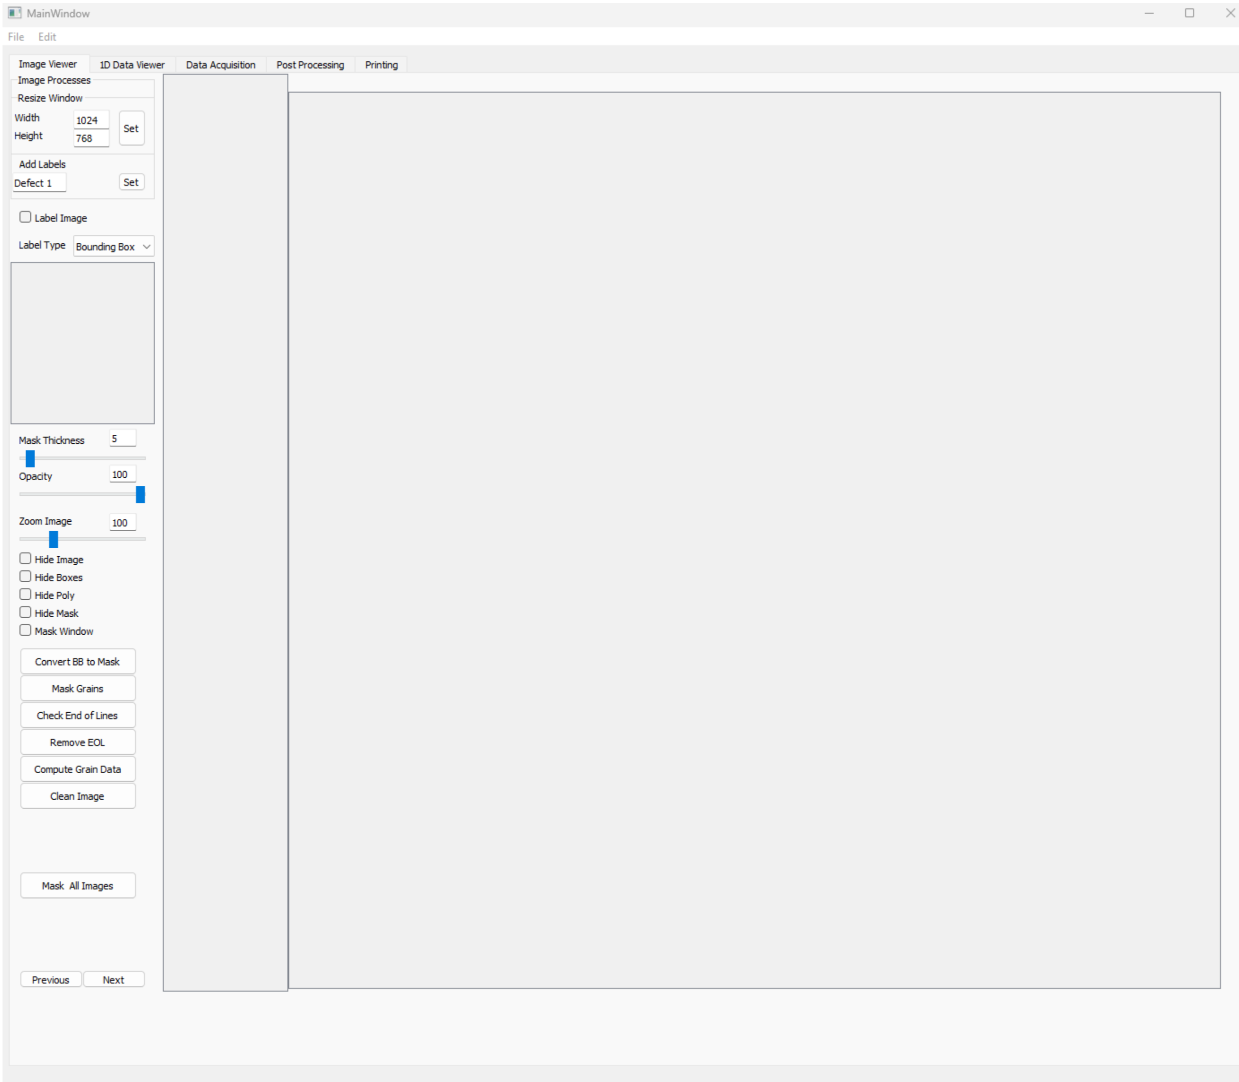

Figure 9.13. SMMVpy2 2D data viewer. Image data can be viewed, and label classes can be added for manual labeling

9.2 XCT Parts

This appendix contains the list of parts used in this project. They contain the XCT results with the pores highlighted in Figure 9.14-Figure 9.25.

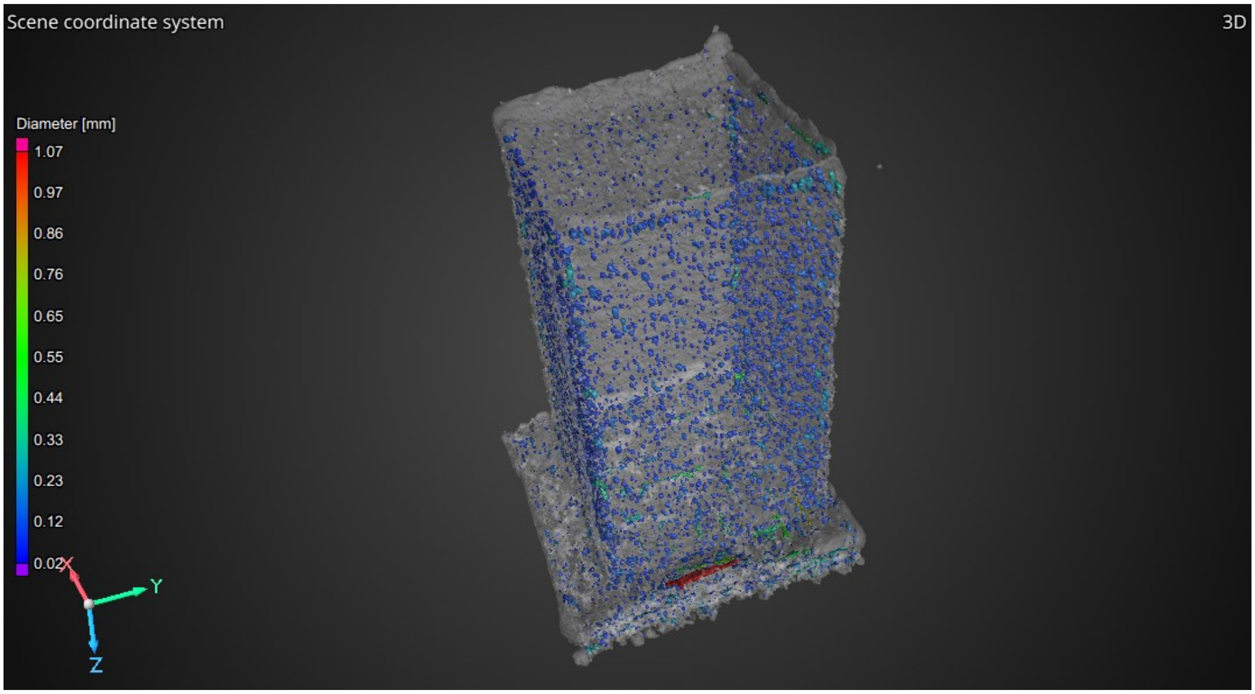

Figure 9.14. XCT result and pores for print 1 part 1.

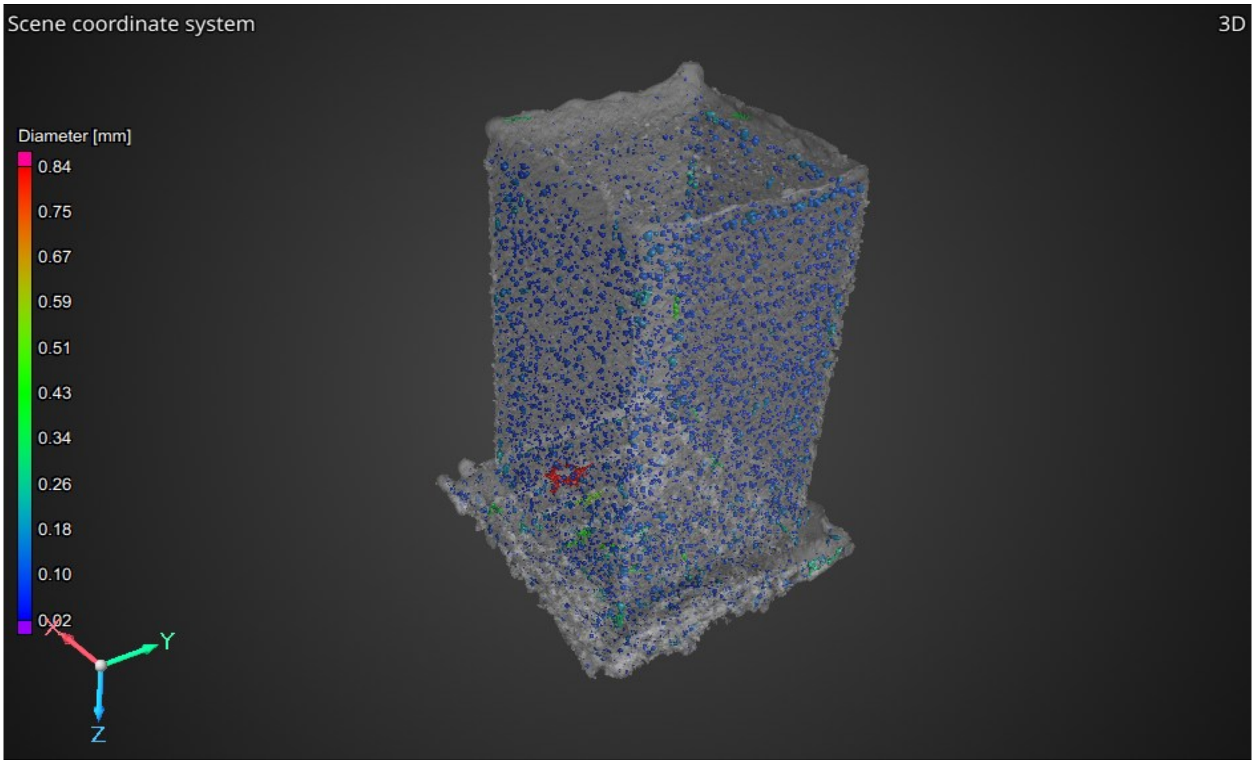

Figure 9.15. XCT result and pores for print 1 part 2.

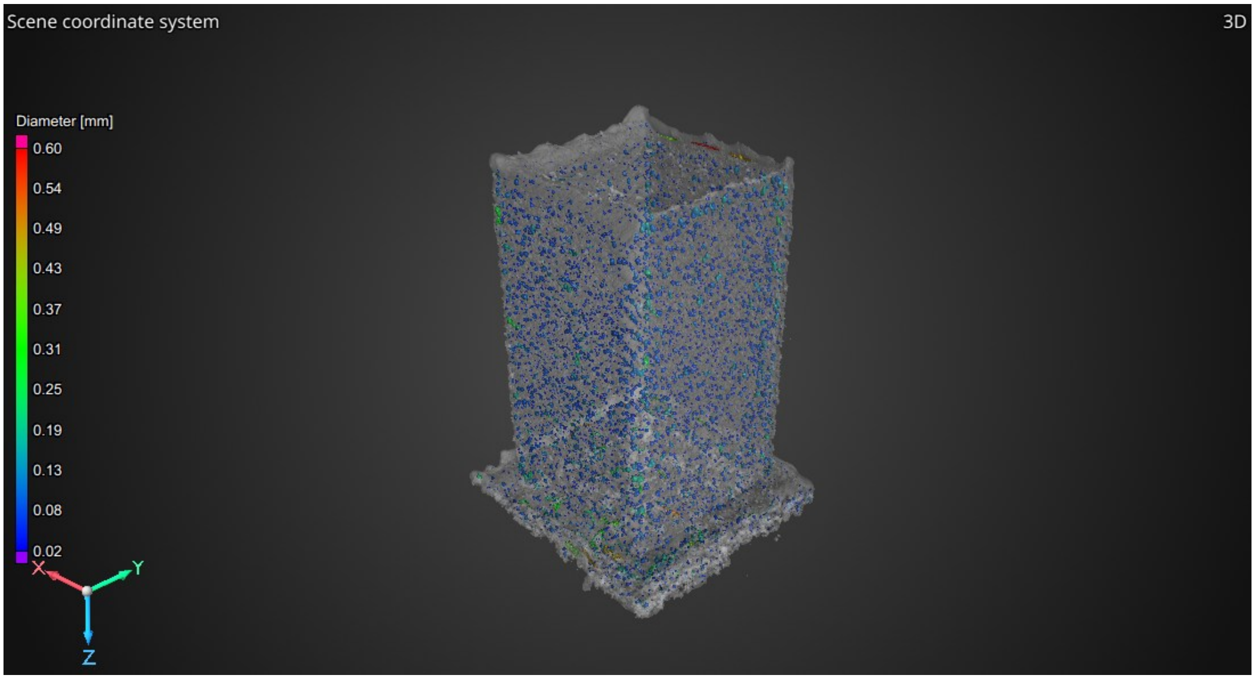

Figure 9.16. XCT result and pores for print 1 part 3.

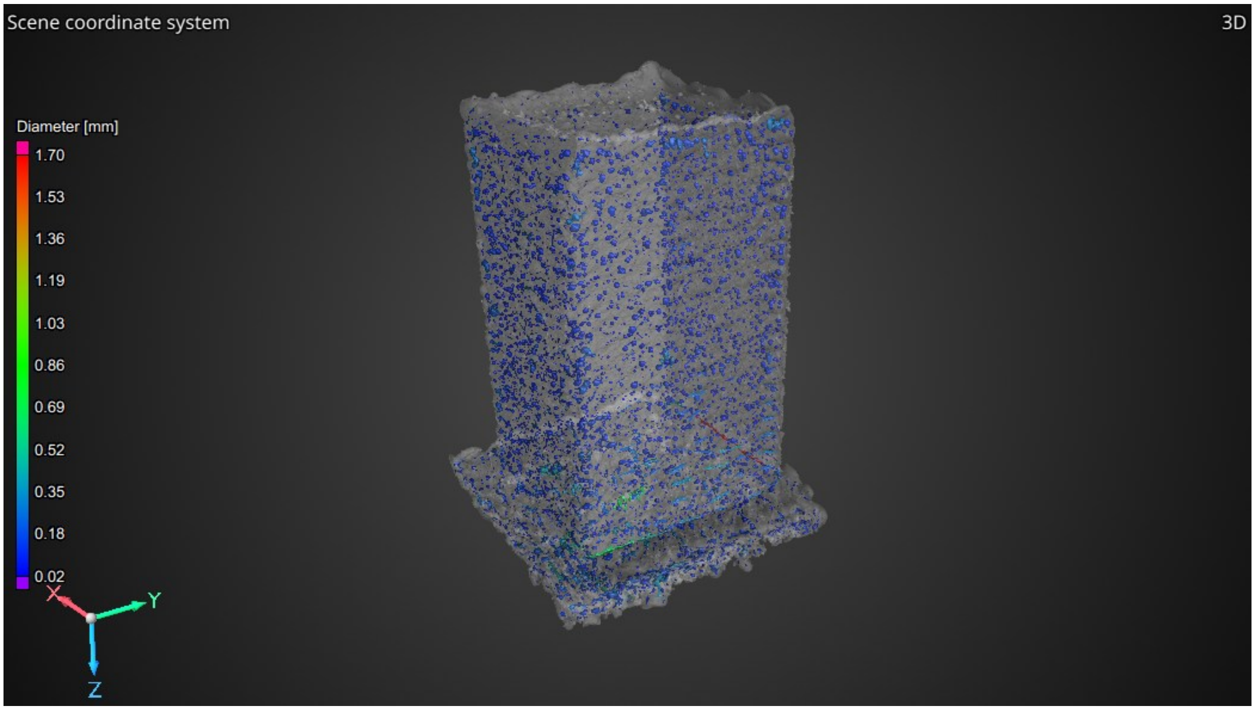

Figure 9.17. XCT result and pores for print 1 part 4.

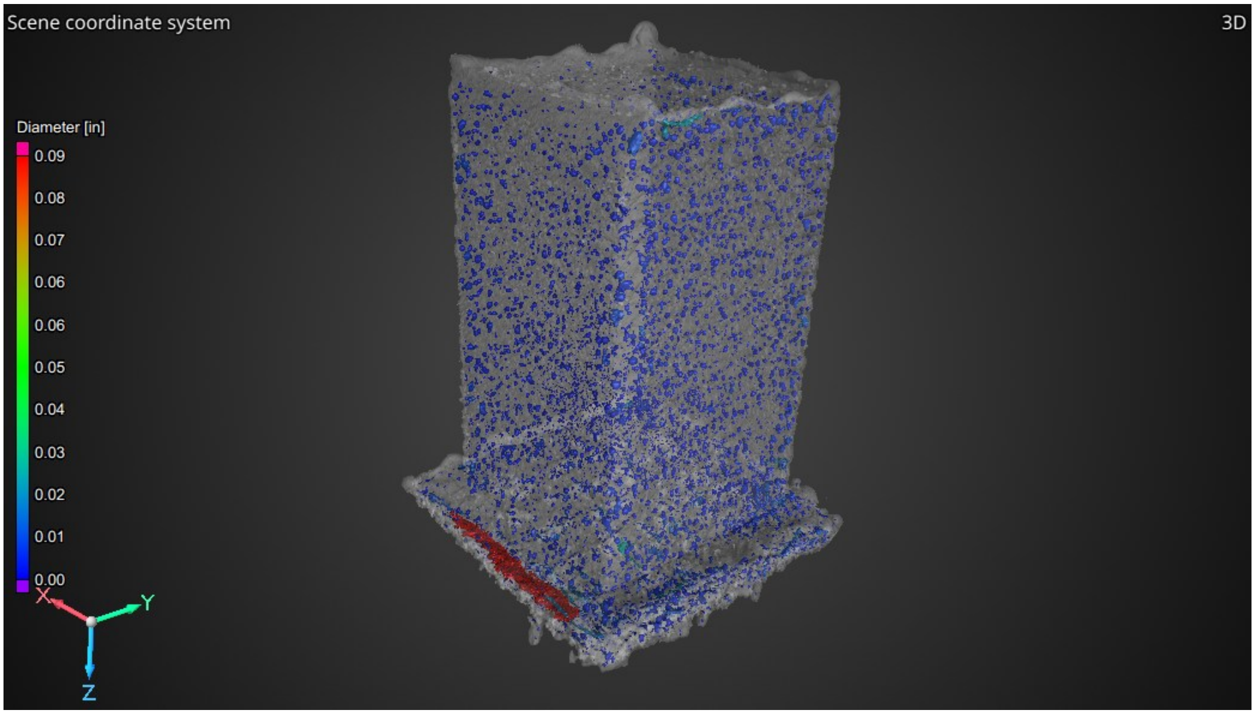

Figure 9.18. XCT result and pores for print 1 part 5.

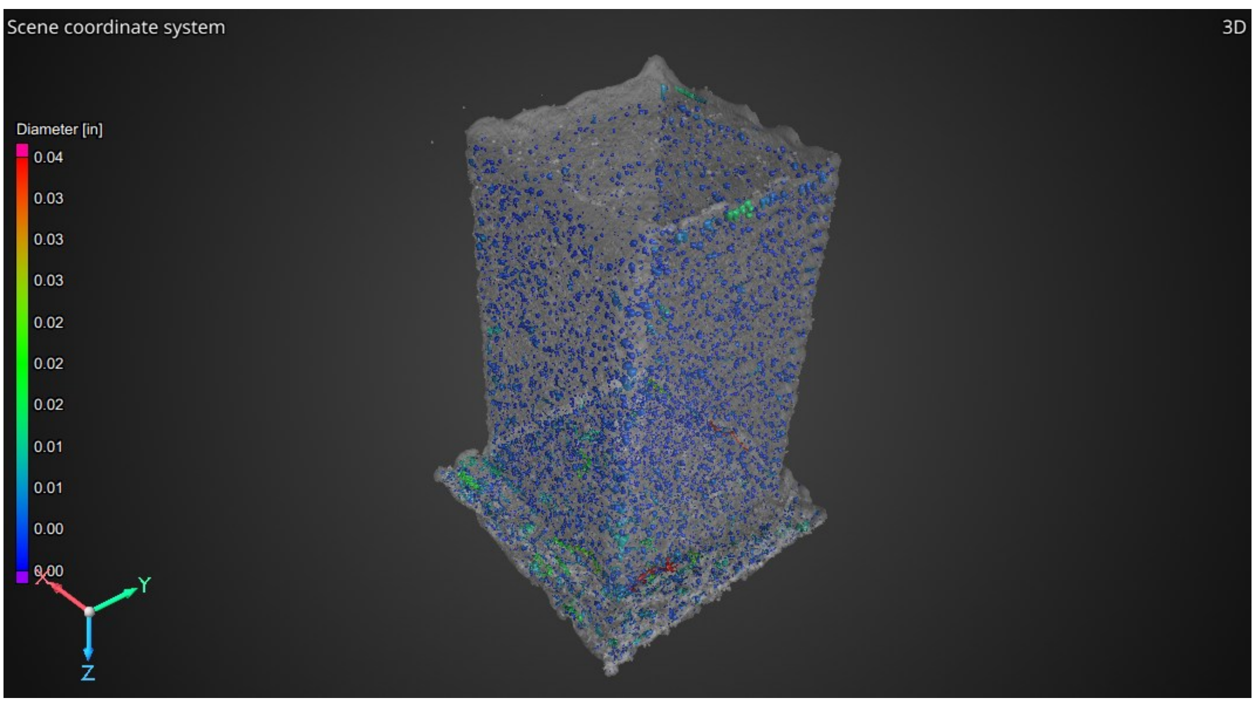

Figure 9.19. XCT result and pores for print 1 part 6.

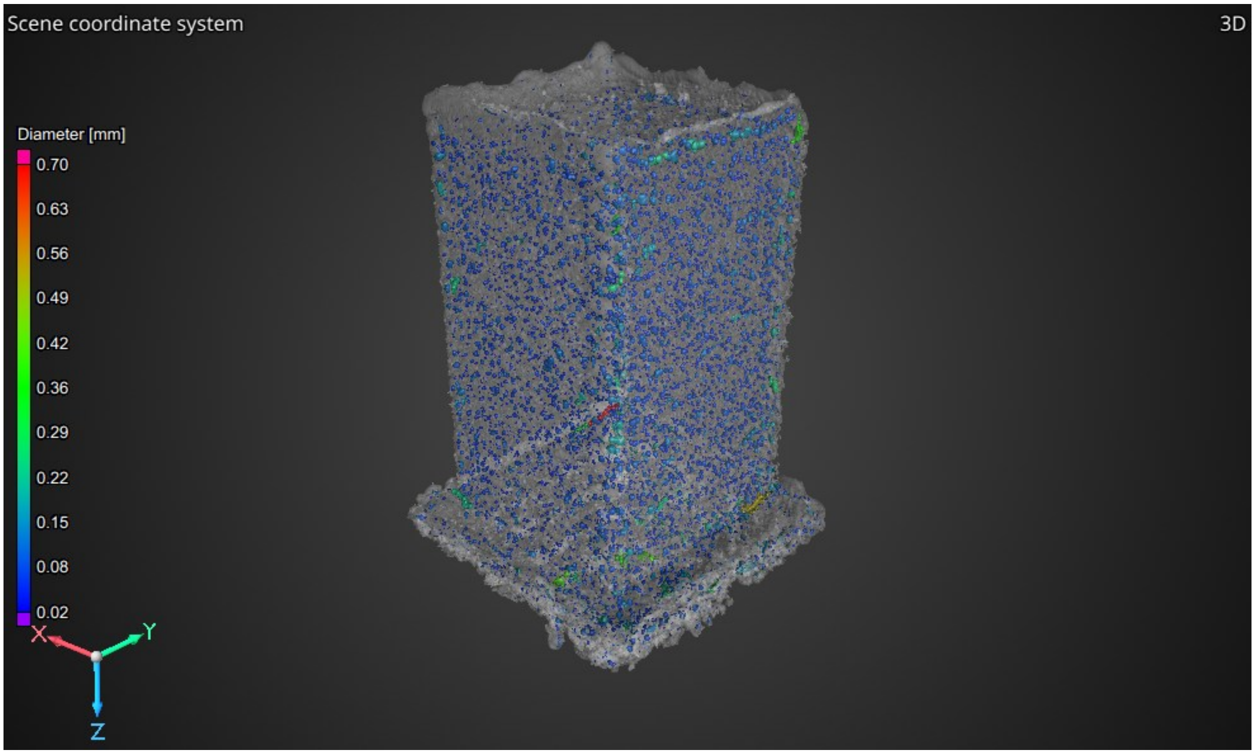

Figure 9.20. XCT result and pores for print 2 part 7.

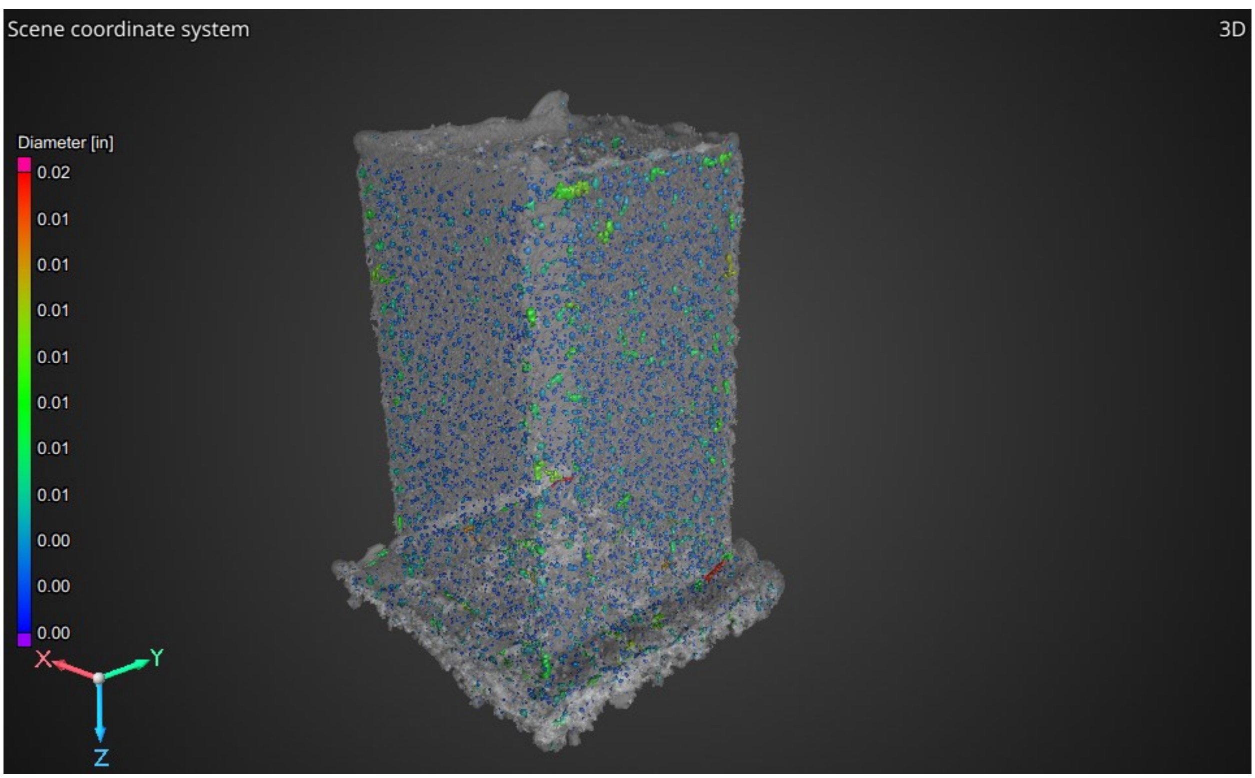

Figure 9.21. XCT result and pores for print 2 part 8.

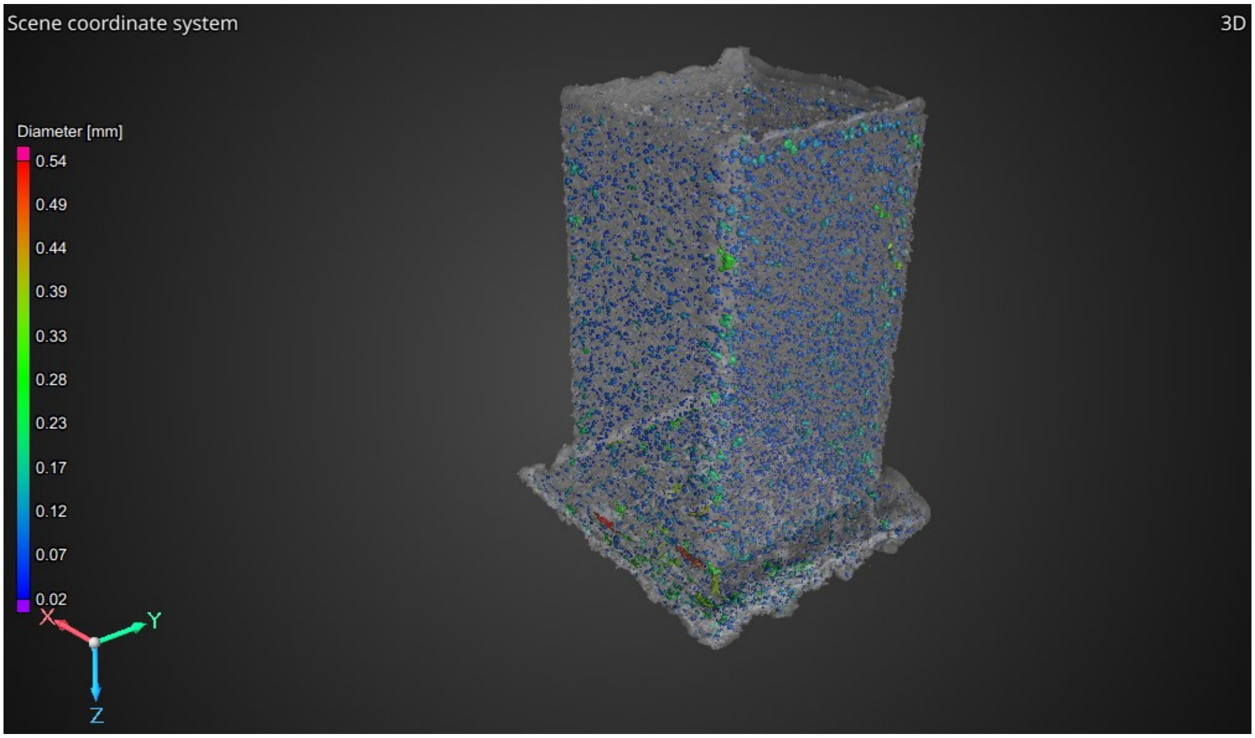

Figure 9.22. XCT result and pores for print 2 part 9.

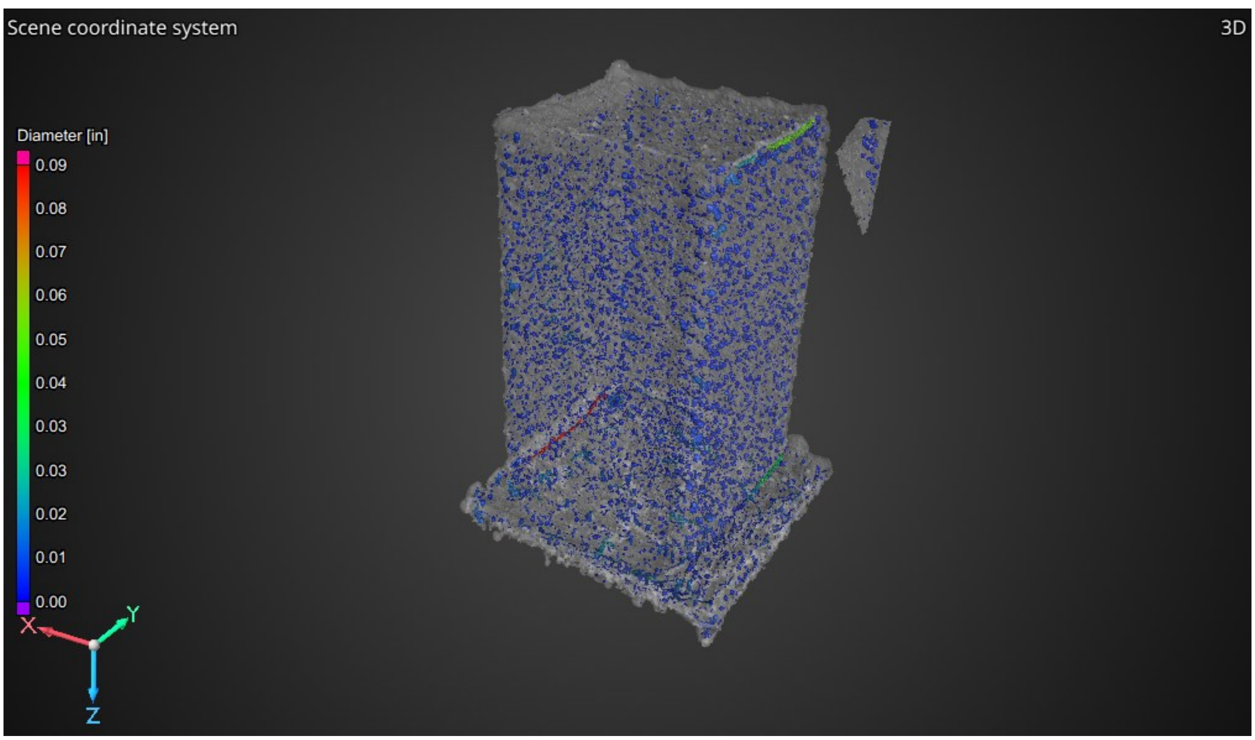

Figure 9.23. XCT result and pores for print 2 part 10. Note that cutoff part is manually removed

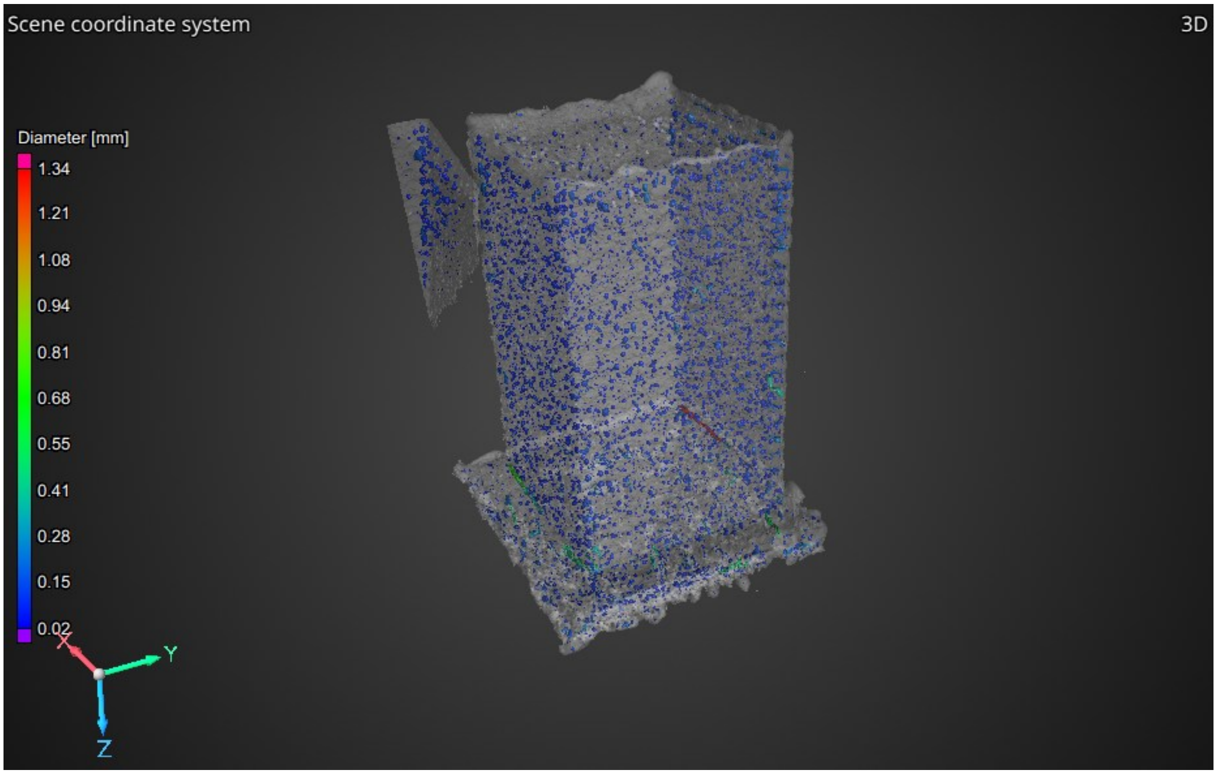

Figure 9.24. XCT result and pores for print 2 part 11. Note that cutoff part is manually removed

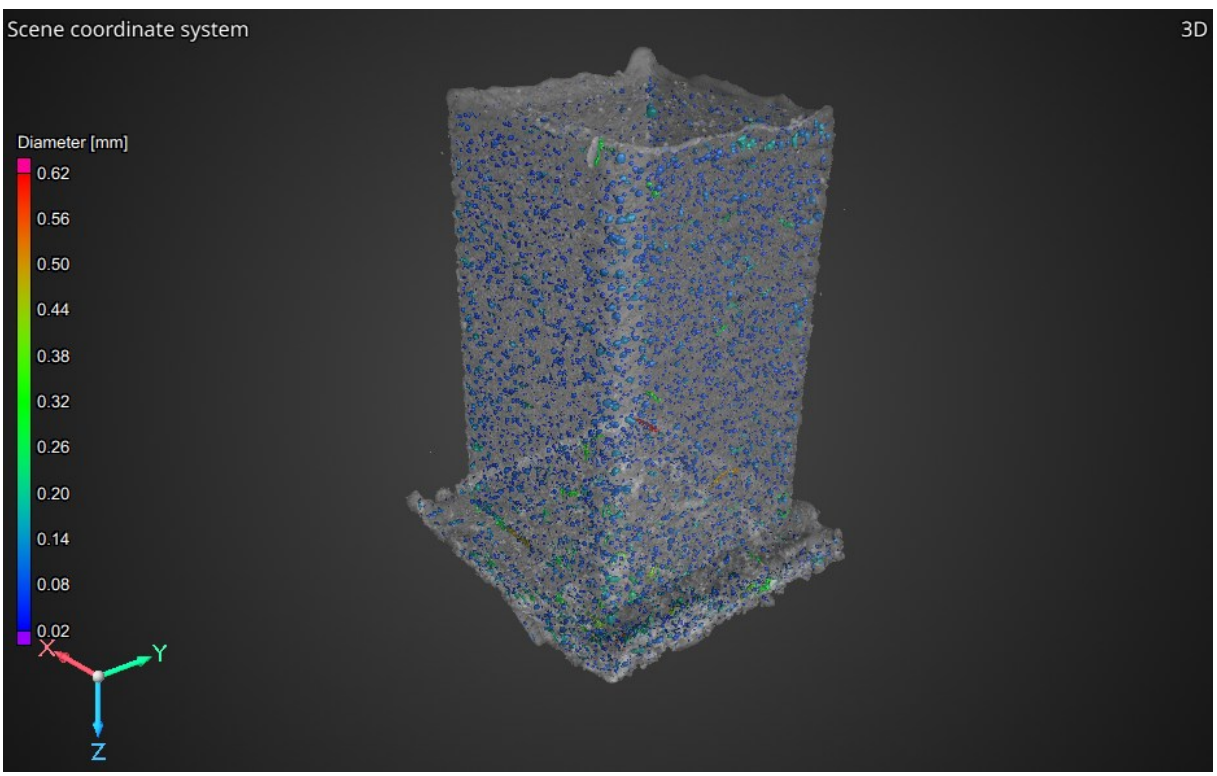

Figure 9.25. XCT result and pores for print 2 part 12.

REFERENCES

Baltrusaitis, Tadas, Chaitanya Ahuja, and Louis-Philippe Morency. “Multimodal Machine Learning: A Survey and Taxonomy.” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 41, no. 2, 2019, pp. 423–443. doi:10.1109/TPAMI.2018.2798607.

Baumgartl, Hermann, Josef Tomas, Ricardo Buettner, and Markus Merkel. “A Deep Learning-Based Model for Defect Detection in Laser-Powder Bed Fusion Using In-Situ Thermographic Monitoring.” Progress in Additive Manufacturing, vol. 5, 2020, pp. 277–285. doi:10.1007/s40964-019-00108-3.

Bidare, P., I. Bitharas, R. M. Ward, M. M. Attallah, and A. J. Moore. “Fluid and Particle Dynamics in Laser Powder Bed Fusion.” Acta Materialia, vol. 142, 2018, pp. 107–120. doi:10.1016/j.actamat.2017.09.051.

Clijsters, Stijn, Tom Craeghs, Sam Buls, Karolien Kempen, and Jean-Pierre Kruth. “In Situ Quality Control of the Selective Laser Melting Process Using a High-Speed, Real-Time Melt Pool Monitoring System.” The International Journal of Advanced Manufacturing Technology, vol. 75, 2014, pp. 1089–1101. doi:10.1007/s00170-014-6214-8.

Cunningham, Ross, Cang Zhao, Niranjan Parab, Christopher Kantzos, Joseph Pauza, Kamel Fezzaa, Tao Sun, and Anthony D. Rollett. “Keyhole Threshold and Morphology in Laser Melting Revealed by Ultrahigh-Speed X-Ray Imaging.” Science, vol. 363, no. 6429, 2019, pp. 849–852. doi:10.1126/science.aav4687.

Davis, Jesse, and Mark Goadrich. “The Relationship Between Precision-Recall and ROC Curves.” Proceedings of the 23rd International Conference on Machine Learning, 2006, pp. 233–240. doi:10.1145/1143844.1143874.

Debroy, T., H. L. Wei, J. S. Zuback, T. Mukherjee, J. W. Elmer, J. O. Milewski, A. M. Beese, A. Wilson-Heid, A. De, and W. Zhang. “Additive Manufacturing of Metallic Components – Process, Structure and Properties.” Progress in Materials Science, vol. 92, 2018, pp. 112–224. doi:10.1016/j.pmatsci.2017.10.001.

Diehl, Brett G., A. Valles Castro, Lars Jaquemetton, and Darren Beckett. “Thermal Calibration of Ratiometric, On-Axis Melt Pool Monitoring Photodetector System Using Tungsten Strip Lamp.” Materials Evaluation, vol. 80, no. 4, 2022, pp. 64–73.

Everton, Sarah K., Matthias Hirsch, Petros Stavroulakis, Richard K. Leach, and Adam T. Clare. “Review of In-Situ Process Monitoring and In-Situ Metrology for Metal Additive Manufacturing.” Materials & Design, vol. 95, 2016, pp. 431–445. doi:10.1016/j.matdes.2016.01.099.

Fawcett, Tom. “An Introduction to ROC Analysis.” Pattern Recognition Letters, vol. 27, no. 8, 2006, pp. 861–874. doi:10.1016/j.patrec.2005.10.010.

Frankot, Robert T., and Rama Chellappa. “A Method for Enforcing Integrability in Shape from Shading Algorithms.” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 10, no. 4, 1988, pp. 439–451. doi:10.1109/34.3909.

Frazier, William E. “Metal Additive Manufacturing: A Review.” Journal of Materials Engineering and Performance, vol. 23, 2014, pp. 1917–1928. doi:10.1007/s11665-014-0958-z.

Galarza, Jose, Jose Barron Jr., Luis Jimenez, Tamer Oraby, Jianzhi Li, and Farid Ahmed. “A Machine Learning Approach to Detect Pores in Laser Powder Bed Fusion Additive Manufacturing.” Manufacturing Letters, vol. 44, Supplement, 2025, pp. 985–993. doi:10.1016/j.mfglet.2025.06.117.

Giri, Shridhar, Sen Liu, Sanam Gorgannejad, Vivek Thampy, Peiyu Quan, Jenny W. Nicolino, Maria Strantza, Jean-Baptiste Forien, Aiden A. Martin, Nicholas P. Calta, and Christopher J. Tassone. “In-Situ Sensor Monitoring of Multi-Class Gas Porosity Formation in Laser Powder Bed Fusion Using Convolutional Neural Network.” Journal of Intelligent Manufacturing, 2026. doi:10.1007/s10845-026-02861-z.

Gorgannejad, Sanam, Aiden A. Martin, Jenny W. Nicolino, Maria Strantza, Gabriel M. Guss, Saad Khairallah, Jean-Baptiste Forien, Vivek Thampy, Sen Liu, Peiyu Quan, Christopher J. Tassone, and Nicholas P. Calta. “Localized Keyhole Pore Prediction during Laser Powder Bed Fusion via Multimodal Process Monitoring and X-Ray Radiography.” Additive Manufacturing, vol. 78, 2023, article 103810. doi:10.1016/j.addma.2023.103810.

Grasso, Marco, and Bianca M. Colosimo. “Process Defects and In Situ Monitoring Methods in Metal Powder Bed Fusion: A Review.” Measurement Science and Technology, vol. 28, no. 4, 2017, article 044005. doi:10.1088/1361-6501/aa5c4f.

Hocker, Samuel J. A., Joseph N. Zalameda, Brodan Richter, Peter W. Spaeth, Andrew R. Kitahara, and Edward H. Glaessgen. “In-Situ Process Monitoring, Synchronization, and Mapping Laser Powder Bed Fusion Builds of Ti6Al4V.” NASA Technical Reports Server, Document ID 20220014500, 2022.

Höfflin, Dennis, Christian Sauer, Alexander Schiffler, Alexander Versch, and Jurgen Hartmann. “Dual Scan Head Approach for In-Situ Defect Detection in Laser Powder Bed Fusion of Metal.” Progress in Additive Manufacturing, vol. 11, 2026, pp. 4077–4096. doi:10.1007/s40964-026-01568-0.

Hofmann, Michael, Thomas Mayer, Francesca Friso, Rene Radis, Christian Kontermann, Falk Müller, and Matthias Oechsner. “Melt Pool Geometry and Process Windows in PBF of 316L: Comprehensive Single-Source Dataset and Statistical Modeling.” Materials & Design, vol. 262, 2026, article 115459. doi:10.1016/j.matdes.2026.115459.

Hooper, Paul A. “Melt Pool Temperature and Cooling Rates in Laser Powder Bed Fusion.” Additive Manufacturing, vol. 22, 2018, pp. 548–559. doi:10.1016/j.addma.2018.05.032.

Huang, Yuze, Tristan G. Fleming, Samuel J. Clark, Sebastian Marussi, Kamel Fezzaa, Jeyan Thiyagalingam, Chu Lun Alex Leung, and Peter D. Lee. “Keyhole Fluctuation and Pore

Formation Mechanisms during Laser Powder Bed Fusion Additive Manufacturing.” Nature Communications, vol. 13, 2022, article 1170. doi:10.1038/s41467-022-28694-x.

Karniadakis, George Em, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. “Physics-Informed Machine Learning.” Nature Reviews Physics, vol. 3, no. 6, 2021, pp. 422–440. doi:10.1038/s42254-021-00314-5.

Khairallah, Saad A., Andrew T. Anderson, Alexander Rubenchik, and Wayne E. King. “Laser Powder-Bed Fusion Additive Manufacturing: Physics of Complex Melt Flow and Formation Mechanisms of Pores, Spatter, and Denudation Zones.” Acta Materialia, vol. 108, 2016, pp. 36–45. doi:10.1016/j.actamat.2016.02.014.

Kim, Felix H., Ho Yeung, and Edward J. Garboczi. “Characterizing the Effects of Laser Control in Laser Powder Bed Fusion on Near-Surface Pore Formation via Combined Analysis of In-Situ Melt Pool Monitoring and X-Ray Computed Tomography.” Additive Manufacturing, vol. 48, 2021, article 102372. doi:10.1016/j.addma.2021.102372.

King, W. E., A. T. Anderson, R. M. Ferencz, N. E. Hodge, C. Kamath, S. A. Khairallah, and A. M. Rubenchik. “Laser Powder Bed Fusion Additive Manufacturing of Metals; Physics, Computational, and Materials Challenges.” Applied Physics Reviews, vol. 2, no. 4, 2015, article 041303. doi:10.1063/1.4936577.

Lane, Brandon M., Shawn P. Moylan, Eric P. Whitenton, and Li Ma. “Thermographic Measurements of the Commercial Laser Powder Bed Fusion Process at NIST.” Rapid Prototyping Journal, vol. 22, no. 5, 2016, pp. 778–787. doi:10.1108/RPJ-11-2015-0161.

Martin, Aiden A., Nicholas P. Calta, Saad A. Khairallah, Jenny Wang, Phillip J. Depond, Anthony Y. Fong, Vivek Thampy, Gabe M. Guss, Andrew M. Kiss, Kevin H. Stone, Christopher J. Tassone, Johanna Nelson Weker, Michael F. Toney, Tony van Buuren, and Manyalibo J. Matthews. “Dynamics of Pore Formation during Laser Powder Bed Fusion Additive Manufacturing.” Nature Communications, vol. 10, 2019, article 1987. doi:10.1038/s41467-019-10009-2.

Oster, Simon, P. P. Breese, A. Ulbricht, G. Mohr, and S. J. Altenburg. “A Deep Learning Framework for Defect Prediction Based on Thermographic In-Situ Monitoring in Laser Powder Bed Fusion.” Journal of Intelligent Manufacturing, vol. 35, 2024, pp. 1687–1706. doi:10.1007/s10845-023-02117-0.

PrintRite3D System Configuration. Manufacturer Delay Settings for PrintRite3D Acquisition System, secToDelayXY = 0.000208 s and secToDelayLDS = 0.000103 s. Configuration used in the synchronized LPBF process-monitoring workflow.

Satterlee, Nicholas, Elisa Torresani, Eugene Olevsky, and John S. Kang. “Automatic Detection and Characterization of Porosities in Cross-Section Images of Metal Parts Produced by Binder Jetting Using Machine Learning and Image Augmentation.” Journal of Intelligent Manufacturing, vol. 35, 2024, pp. 1281–1303. doi:10.1007/s10845-023-02100-9.

Satterlee, Nicholas, Runjian Jiang, Eugene Olevsky, Elisa Torresani, Xiaowei Zuo, and John S. Kang. “Robust Image-Based Cross-Sectional Grain Boundary Detection and Characterization Using Machine Learning.” Journal of Intelligent Manufacturing, vol. 36, 2025, pp. 3067–3095. doi:10.1007/s10845-024-02383-6.

Satterlee, Nicholas, Xiaowei Zuo, Chang-Whan Lee, Choon-Wook Park, and John S. Kang. “Parallel Multi-Layer Sensor Fusion for Pipe Leak Detection Using Multi-Sensors and Machine Learning.” Engineering Applications of Artificial Intelligence, vol. 153, 2025, article 110923. doi:10.1016/j.engappai.2025.110923.

Scime, Luke, and Jack Beuth. “A Multi-Scale Convolutional Neural Network for Autonomous Anomaly Detection and Classification in a Laser Powder Bed Fusion Additive Manufacturing Process.” Additive Manufacturing, vol. 24, 2018, pp. 273–286. doi:10.1016/j.addma.2018.09.034.

Shevchik, Sergey A., Christoph Kenel, Christian Leinenbach, and Kilian Wasmer. “Acoustic Emission for In Situ Quality Monitoring in Additive Manufacturing Using Spectral Convolutional Neural Networks.” Additive Manufacturing, vol. 21, 2018, pp. 598–604. doi:10.1016/j.addma.2017.11.012.

Tang, Q., Liang, J., Zhu, F., 2023. A comparative review on multi-modal sensors fusion based on deep learning. Signal Process. 213, 109165. https://doi.org/10.1016/j.sigpro.2023.109165.

Tapia, Gustavo, and Alaa Elwany. “A Review on Process Monitoring and Control in Metal-Based Additive Manufacturing.” Journal of Manufacturing Science and Engineering, vol. 136, no. 6, 2014, article 060801. doi:10.1115/1.4028540.

Thompson, Adam, Ian Maskery, and Richard K. Leach. “X-Ray Computed Tomography for Additive Manufacturing: A Review.” Measurement Science and Technology, vol. 27, no. 7, 2016, article 072001. doi:10.1088/0957-0233/27/7/072001.

Woodham, Robert J. “Photometric Method for Determining Surface Orientation from Multiple Images.” Optical Engineering, vol. 19, no. 1, 1980, pp. 139–144. doi:10.1117/12.7972479.

Yeung, Ho, and Steven Grantham. “Laser Calibration for Powder Bed Fusion Additive Manufacturing Process.” Solid Freeform Fabrication Symposium 2022, Austin, TX, 2022. National Institute of Standards and Technology, https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=935350.

Zhao, Cang, Niranjan D. Parab, Xuxiao Li, Kamel Fezzaa, Wenda Tan, Anthony D. Rollett, and Tao Sun. “Critical Instability at Moving Keyhole Tip Generates Porosity in Laser Melting.” Science, vol. 370, no. 6520, 2020, pp. 1080–1086. doi:10.1126/science.abd1587.